

UC Berkeley

UC Berkeley Previously Published Works

Title

Building the Second Mind: 1956 and the Origins of Artificial Intelligence Computing

Permalink

<https://escholarship.org/uc/item/88q1j6z3>

ISBN

9781476357638

Author

Skinner, Rebecca Elizabeth

Publication Date

2012-05-01

Building the Second Mind: 1956 and the Origins of Artificial Intelligence Computing

Rebecca E. Skinner

copyright 2012 by Rebecca E. Skinner
ISBN 978-0-9894543-1-5

Building the Second Mind: 1956 and the Origins of Artificial Intelligence Computing

Chapter .5. Preface

Chapter 1. Introduction

Chapter 2. The Ether of Ideas in the Thirties and the War Years

Chapter 3. The New World and the New Generation in the Forties

Chapter 4. The Practice of Physical Automata

Chapter 5. Von Neumann, Turing, and Abstract Automata

Chapter 6. Chess, Checkers, and Games

Chapter 7. The Impasse at End of Cybernetic Interlude

Chapter 8. Newell, Shaw, and Simon at the Rand Corporation

Chapter 9. The Declaration of AI at the Dartmouth Conference

Chapter 10. The Inexorable Path of Newell and Simon

Chapter 11. McCarthy and Minsky begin Research at MIT

Chapter 12. AI's Detractors Before its Time

Chapter 13. Conclusion: Another Pregnant Pause

Chapter 14. Acknowledgements

Chapter 15. Bibliography

Chapter 16. Endnotes

Chapter .5. Preface

Introduction

Building the Second Mind: 1956 and the Origins of Artificial Intelligence Computing is a history of the origins of AI. AI, the field that seeks to do things that would be considered intelligent if a human being did them, is a universal of human thought, developed over centuries. Various efforts to carry this out appear in the forms of robotic machinery and more abstract tools and systems of symbols intended to artificially contrive knowledge. The latter sounds like alchemy, and in a sense it certainly is. There is no gold more precious than knowledge. That this is a constant historical dream, deeply rooted in the human experience, is not in doubt. However, it was not more than a dream until the machinery that could put it into effect was relatively cheap, robust, and available for ongoing experimentation.

The digital computer was invented during the years leading to and including the Second World War, and AI became a tangible possibility. Software that used symbols to enact the steps of problem-solving could be designed and executed. However, envisioning our possibilities when they are in front of us is often a more formidable challenge than bringing about their material reality. AI in the general sense of intelligence cultivated through computing had also been discussed with increasing confidence through the early 1950s. As we will see, bringing it into reality as a concept took repeated hints, fits, and starts until it finally appeared as such in 1956.

Our story is an intellectual saga with several supporting leads, a large peripheral cast, and the giant sweep of Postwar history in the backdrop. There is no single 'great man' in this opus. As far as the foundation of AI is concerned, all of the founders were great. Even the peripheral cast was composed of people who were major figures in other fields. Nor, frankly, is there a villain either.

Themes and Thesis

The book tells the story of the development of the cognitive approach to psychology, computer science (software), and the development of software that undertook to do 'intelligent' things

during mid-century. To this end, I study the early development of computing and psychology in the middle decades of the century, ideas about 'Giant Brains', and the formation of the field of study known as AI.

Why did this particular culture spring out of this petri dish, at this time? In addition to 'why', I consider the accompanying where, how, and who. This work is expository: I am concerned with the enrichment of the historical record. Notwithstanding the focus on the story, the author of necessity participates in the thematic concerns of historians of computer science. Several themes draw our attention.

The role of the military in the initial birth and later development of the computer and its ancillary technologies should not be erased, eroded, or diminished. Make no mistake: war is abhorrent. But sustained nation-building and military drives can yield staggering technological advances. War is a powerful driver of technological innovations (1). This is particularly the case with the development of 'general-purpose technologies', that is, those which present an entirely new way of processing material or information (2). These technologies of necessity create and destroy new industries, means of locomotion, creation of energy, and processing of information (steel rather than iron, the book, the electric generator, the automobile, the digital computer). In the process, these fundamental technologies will bring about new forms of communication, cultural activities, and numerous ancillary industries. We repeat, for effect: AI is the progeny of the Second World War, as is the digital computer, the microwave oven, the transistor radio and portable music devices, desktop and laptop computers, cellular telephones, the iPod, iPad, computer graphics, and thousands of software applications. The theory of the Cold War's creative power and fell hand in shaping the Computer Revolution is prevalent in the current academic discourse on this topic: its paradoxical creative power can't be denied (3).

The role of the Counterculture in creating the Computer Revolution is affectively appealing. In its strongest form, this

theory holds that revolutionary hackers created most, if not all of the astonishing inventions in computer applications (4). For many of the computer applications of the Sixties and Seventies, including games, software systems, security and vision, this theory holds a good deal of force. However, the period under discussion in this book refutes the larger statement, though. The thesis has its chronology backwards. The appearance of the culturally revolutionary 't-shirts' was preceded by a decade and a half of hardware, systems and software language work by the culturally conservative 'white-shirts'. (Throughout the Fifties, IBM insisted on a dress code of white shirts, worn by white Protestant men) (5).

Yet there is one way in which those who lean heavily on the cultural aspect of the Computer Revolution are absolutely correct. An appropriate and encouraging organizational culture was also essential in the development of the computer in every aspect, along the course of the entire chronology. This study emphasizes more emphatically the odd *mélange* of a number of different institutional contexts in computing, and how they came together to study one general endeavor. AI in its origins started with individual insights and projects, rather than cultivated in any single research laboratory. The establishment of AI preceded its social construction. We could say that AI's initial phase as revolutionary science (or exogenous shock, in the economist's terms) preceded its institutionalization and establishment of an overall "ecology of knowledge" (6). However, once AI was started, it too relied heavily on its institutional settings. In turn, the founders of AI established and cultivated research environments that would continue to foster innovation. The cultivation of such an environment is more evident in the later history of AI, rather than in the tentative movements in the 1950s.

Yet another salient theme of this work is the sheer audacity of the paradigm transition that AI itself entailed. The larger manner of thinking about intelligence as a highly tangible quality, and about thinking as something that could have qualitative aspects to it, required a vast change between the late 1930s and the mid-1950s. As with any such alteration of focus, this required the new

agenda to be made visible- and envisioned while it still seemed like an extreme and far-fetched concept (7).

A final theme is the larger role of AI in the history of the Twentieth century. It is certainly true that the Cold War's scientific and research environment was a 'Closed World' in which an intense, intellectually charged, politically obsessive culture thrived. The stakes involved in the Cold War itself were the highest possible ones; the intensity of its inner circles this is no surprise. However, in this case, the larger cultural themes of the Twentieth century had created their own "closed world". Between them, Marxian political philosophy and Freudian influence on culture had robbed the arts, politics and literature of its vitality. This elite literary 'Closed world' saw science and technology as aesthetically unappealing and inevitably hijacked by the political forces that funded research. The resolution of the Cold War, and the transformation of the economy and ultimately of culture by the popularization of computing, would not take place for decades. Yet the overcoming of the cultural impasse of the Twentieth century would be a long-latent theme in which AI and computing would later play a part (8).

Outline of the Text

In Building the Second Mind: 1956 and the Origins of Artificial Intelligence Computing, we examine the way in which AI was formed at its start, its originators, the world they lived in and how they chose this unique path, the computers that they used and the larger cultural beliefs about those machines, and the context in which they managed to find both the will and the way to achieve this. 1956 was the tipping point, rather than the turning point, for this entry into an alternative way of seeing the world. Our chapter outline indicates the line the book follows.

The chapter outline delineates the book's course. The Introduction and Conclusion chapters frame the book and discuss history outside of the time frame covered in BTSM, and establishes AI as a constant in world intellectual history (Chapter One). The other chapters are narrative, historical, and written within the context of their time.

The intellectual environment of the Thirties and the war years is foreign to us. It lacked computers, and likewise lacked any sort of computational metaphor for intelligence. Studies of intelligence without any real reference to psychology nevertheless abounded (Chapter Two). Cognitive psychology was not a topic of academic study through the first two quarters of the Twentieth century. Yet intelligent processes and learning were discussed in multifold ways. Formal logic developed symbolic languages for the representation of the real world by a symbolic language. This language was not yet computational. However, information theory, which concerned the integrity of transmission of electrical signals, was invented by Claude Shannon, Warren Weaver, and other engineers. The latter two things had not yet been joined in the implementation of computer languages- but this was a development that could have been predicted.

Revisiting the late 1940s, one is struck by the sheer foreignness of the environment (Chapter Three). The political environment, with the ominous Cold War between the Soviet Union and the United States, and practically every other country damaged by the Second World War, seems firmly in another century. The financial and academic anemia of the 1930s gave way to the wartime research programs of the 1940s, and brought with it opportunities for a new generation. The universities expanded, and many new research institutions were established. Moreover, the ongoing development of the computer and other technologies that benefitted from the Cold War offered opportunities unimaginable before 1945. The generation that benefitted from these new circumstances, too young to have served in the Second World War, or to have had their personal histories and careers marred by the Depression, was indeed fortunate in the timing of their lives. The leaders of AI for its first decades- John McCarthy, Marvin Lee Minsky, Allen Newell, and Herbert Alexander Simon- were uniquely favored by history. These four, and their larger cohort, benefitted as well from the increasingly open intellectual environment of the Cold War, as professional societies, the universities, research institutes, and

even popular media all encouraged the discussion of computing, intelligence and its emulation (Chapter Four).

Continually repressed by the state of academic psychology, the study of intelligence further made its appearance in the design of other intelligent beings besides robotic artifacts (Chapter Five). Singular minds such as John Von Neumann and Alan Turing proposed intelligent automata, which would be binary coded programs that could carry out computable functions (i.e., equations). Von Neumann's discussion of automata, and the work of Warren McCulloch and Walter Pitts, suggested that these proposed creations be made of biological matter- essentially A-Life before its time. Turing also designed the eponymous Turing Test, a means of determining whether a given intelligent machine was actually intelligent. Both Turing and Von Neumann died before AI had advanced very far; however they greatly influenced their generation in general and Minsky and McCarthy in particular.

If electric robotic automata were one prevalent form of the emulation of intelligence throughout the 20th century, games were another form and often represented the higher ground for such representation (Chapter Six). Chess is too large a search space for undirected 'blind' search, so it immediately challenged the early users of computers toward strategy. Claude Shannon used the gedankenexperiment of chess as a backdrop for the visualization of problem-solving as an upside-down tree the branches and branching points of which can be portrayed as positions and moves respectively. Other computer scientists, at Los Alamos and at the National Bureau of Standards and the ever-busy Newell and Simon and their colleague Clifford Shaw, began to work on games, often as a lark when the workday at the computer was over. Finally, IBM programmer Arthur Samuel developed a checkers-playing computer program that was so popular with the media that IBM sent him to Europe to try to hush the attention.

Early in the Fifties, ideas began to advance far ahead of technological expression. The grand automata envisioned by Von Neumann, Turing, Warren Weaver and others could not be realized practically. Computers worked very poorly; there was no

operating software and all programs had to be written by hand and executed slowly during limited hours. The earliest digital computers were initially characterized in the popular media, as dangerous and preposterous machines that were visually akin to 'giant brains' on legs. There was scarcely any terminology for any aspect of human intelligence. Cognitive psychology and its close studies of conventions in thought processes were finally initiated, but existed on a very small scale. Cybernetics was stalled, dominated by Norbert Wiener's increasingly maudlin statements as to the fear of nuclear war. Technological advances were required to dispel the impasse through ongoing progress on all fronts (Chapter Seven).

In Chapter Eight, we will bring in the essential role of the Rand Corporation in providing an amenable petri dish for some of the earliest AI programs. The Rand Corporation in Santa Monica, the think tank spun off from Air Force research after the Second World War, became the location of choice for Cold Warriors discussing nuclear war strategies during the Fifties and Sixties. Rand boasted a rare digital computer, used for calculations of war games. Herbert Simon, later joined by Allen Newell, began to consult there in 1952. Clifford Shaw, a programmer at Rand, worked with them to develop the Logic Theorist. Using Rand's Johnniac computer, they devised this program, which is given uncompleted logic theorems to prove. Newell and Simon brought this program's evidence to the Dartmouth Conference in 1956.

The Dartmouth Summer Conference, held in 1956 at Dartmouth University in Hanover, New Hampshire, brought together the significant participants in AI during its first two decades (Chapter Nine). This was the tipping point at which AI was established as such; at which the name AI, 'artificial intelligence' was first widely used; and at which the general trends and differences were clarified. It is also the moment at which NSS' prospectively more felicitous term of 'complex information processing' was rejected. However, the clarification of AI as a common effort established the founders as a group with common beliefs about the possibility of AI itself.

The Dartmouth Conference did not change the research orientation of any of its participants, but it did establish the concept and name of AI, and present it to a distinguished cohort of the computer and Cybernetics milieu. Over the next several years, the Cybernetic agenda for intelligence and its biological metaphors for intelligence was prevalent. While it would eventually dwindle simply due to its lack of connection to a scientific research paradigm, its dissipation would take a number of years. The next two chapters examine the progress of AI during the remainder of the 1950s and into the first year of the 1960s. At Carnegie Tech, Newell and Simon continued on their inexorable path, using chess-playing and the ambitiously named General Problem Solver program as a way to establish a vocabulary for cogitation (Chapter Ten). The latter was impressive but certainly not what it claimed to be, but AI has done well when it has aimed for big things, even if those things were not achieved immediately.

The next chapter follows McCarthy and Minsky (Chapter Eleven). Working at Dartmouth, then MIT for several years, then settling finally at Stanford, John McCarthy made foundational contributions to the field in the form of the LISP computer language and the invention of timesharing. At MIT, Marvin Minsky worked at the Research Laboratory of Electronics, and then joined the Mathematics department in 1958. He and McCarthy established the Artificial Intelligence Project the same year, and began to gather an eager undergraduate following of student hackers, who initiated research in visual display, computer games, and graphics.

Like any audacious idea, AI attracted detractors who asserted that it was impossible, lacking in sensitivity toward intelligence itself, and overly audacious (Chapter Twelve). Once AI was actually an extant aspiration, it garnered much more bile and publicity than the detractors desired. During the first several years, the most prominent and persistent of the detractors appeared and began criticizing the field on the grounds that AI could not grasp the phenomenological nature of human perception. This is entirely true of early AI- even truncated

cogitation is difficult to embody- but the extremely negative tone meant that it was not constructive criticism.

As our story ends at the conclusion of the 1950s, AI had built a foothold in the major institutions which would nurture it over the next decades (Chapter Thirteen). Its founders had their research funded, and the early major achievements were extant. The intellectual orientation which demands usage of computer programs to try to explore cognition was well-established. Gradually but surely, this paradigm was supplanting a cybernetic orientation which took its cues from engineering.

Chapter 1. Introduction: The Prehistory of AI

The Conceptual Watershed of the Dartmouth Conference

In the summer of 1956, roughly two dozen men gathered at the bucolic rural campus of Dartmouth University in Hanover, New Hampshire for a conference designated 'the Dartmouth Summer Research Project on Artificial Intelligence.' Dartmouth professor John McCarthy, who had initially suggested the conference, had explained in the conference's original proposal:

" A two-month, ten-man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth in Hanover, N.H. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it" (9).

This was a strange and novel concept at this time- on a par with the proposal of a round Earth, a united nations, universal or female suffrage, the abolition of slavery, the separation of church and state, evolution through genetic mutation and natural selection, or plate tectonics, in their respective times. The possibility of computers simulating thinking was outside the boundaries of their stipulated tasks. These consisted primarily of adding large rows of numbers for ballistic, actuarial, and meteorological equations and the like. The computer of the future

had even been portrayed as a giant brain on legs- a monstrous and weird image with no practical import. The idea of computers that engaged in anthropomorphically appealing, 'intelligent' activities was as immediately probable and appealing to common wisdom as suggesting a portable or household nuclear reactor to fill one's heating and electrical needs. One of the world's foremost computing authorities had stated just a few years earlier that the world market for computers would never exceed a half-dozen (10). In this context, it is easy to see why AI lacked not only the respect of many computer professional, but also obvious, immediate clues as to how to proceed.

Today, computing applications number in the hundreds of thousands, instantaneous wireless connectivity is taken for granted, and computers and PDAs are ubiquitous in the lives of every middle-class person on the planet. But in 1956, the world was unimaginably different- computationally speaking. Computing was slow, ponderous, and involved input of both programs and data through paper cards or even paper tape, and weak and fallible magnetic core memory. The number of computers in the world could almost be counted precisely, because they were all owned by governments or major corporations or research centers (11). Storage took place on magnetic tape (a rare computing and storage medium that persists even today). There were no operating systems, no systems or commercial applications software, no computing application stores, wireless connectivity, tech support (online or on the phone), Apple Stores, Geek Squad, no online world at all.

A plethora of weaknesses, including lack of choice of vendor or end product, inflexibility of use, lack of universality, and slow progress in consumer product development, characterized information technology itself. Clunky, rotary-dialed telephones were attached to the wall by wires. Phone calls, especially long-distance, were expensive and often had to be arranged with a live operator. Poor-quality mimeographs producing smudged purple copies proliferated, as did telegraphy for terse, important, long-distance messages. Even robust electric typewriters did not exist: IBM did not introduce its iconic Selectric model until 1961. A

“computer” was sometimes still understood in its archaic form- as a person, typically a female, who conducted mathematical operations using pencil and paper. Machine-readable (MICR) numbers still had not been invented; penmanship was an important topic in primary education. Handwriting forgery was a serious forensic concern. In an environment in which data flowed over electronic networks with such expense and difficulty, the fluid movement of data in digital form was barely conceivable.

The introduction and actual implementation of an idea that was practically speaking before its time- and which always had been- required a will and a way. The way was present in the reality of the general-purpose digital computer, invented barely a decade earlier and under increasingly intense development during the entire second half of the Twentieth century. The will required more historical serendipity, in the form of several brilliant- but extremely different- men who became the field’s founders.

John McCarthy invented and developed the first widely-used list-processing language, which was essential to the clear expression of AI’s concepts. Marvin Minsky, known for the concept of the society of mind, which sees intelligence as numerous agencies, or independent capacities. He is also the AI founder best-known to the general public, and the most involved with robotics and artificial vision. Allen Newell and Herbert Simon generally approached AI with a concern for its contributions to cognitive psychology. Newell and Simon, working with programmer Cliff Shaw, also produced the earliest AI programs.

Nothing in the academic environment or the world at large suggested that this enterprise would realize any of its goals soon. No matter; the conceptual watershed reached by this group’s simple insistence on inquiry into the idea itself mattered enormously.

This book tells the story of the establishment of AI by a handful of people during the 1950s. Both the larger historical environment and the staggering ingenuity of several people led to this idea

being one that could finally be realized, after centuries as a dream.

The Present in the Past: Origins in AI's Antiquity

“ Like the old woman in the story who described the world as resting on a rock, and then explained that rock to be supported by another rock, and finally when pushed with questions said it was rocks all the way down- he who believes this to be a radically moral universe must hold the moral order to rest either on an absolute and ultimate should, or on a series of shoulds all the way down” (12).

In looking for the origins of AI, it is tempting to say that it has been with humanity forever- or rocks or turtles or elephants “all the way down”. The idea of understanding intelligence systematically and emulating it in machinery- or biology in the form of human like automata- is ancient rather than recent. AI is a new science with an ancient heritage in philosophy and automata. Explaining the timelessness of its origins repudiates the ostensible delusions of its aspirations at that time.

The Ancient Greeks and the Earliest Cognitive Sciences

AI's first conceptual precedents are found in Athenian Greece in the sixth century B.C.E. Attic Greece was not much concerned with numbers per se, but was fertile with other aspects of understanding intelligence and ideas of the mind. This society produced the earliest statements of the reality of abstract ideas; the first geometric proofs and efforts at logical forms for argumentation; and a rudimentary theory of the mind. Plato (429-347 B.C.E.) drew out the idea of Platonic absolutes or absolute forms, of all material objects, in The Dialogues, a series of confrontational conversations with his teacher, Socrates. This was the first effort at knowledge representation, or clear ways to speak about different sorts of ideas. Moreover, Plato's early proofs seem to be efforts at formalized problem-solving. The work of Plato's student, Aristotle (384-322 B.C.E.), features a systematic search for answers to questions of natural science.

This approaches a prehistoric paraphrase of the initial state and the goal state, or what has come to be known in AI as search through all sorts of problem spaces.

If the reality of abstractions is central, so is the necessity of establishing rules or some other format for approaching problem-solving or representation of concepts. The concept of protocols was first touched upon with the Attic Greek concept of heuristics—defined by mid-20th century philosopher George Polya as “an adjective, [which] means serving to discover” (13). Finally, the philosopher Euclid’s representation of geometrical figures in imagined space has been established as part of the field of geometry (14).

Automata, Mechanical and Biological

Automaton: “a mechanism that is relatively self-operating, especially robot; a machine or control mechanism designed to follow automatically a predetermined sequence of operations or respond to encoded instructions.” (Webster’s Ninth New Collegiate Dictionary).

Recreating a human being has been one of history’s most persistent nostalgias. The word is derived from the Latin, in turn based on the Ancient Greek *automatos*, referring to a machine with the ability to move by its own force. Biological automata, are universal. Purported recreation of human body and intelligence is prefigured in the Babylonian Epic of Gilgamesh, and in the Hebrew conceptualization of humans as made by God from the earth. Mechanical automata, and prosthetics with automata-like features, abound in ancient and especially Greek myth and practice. According to legend, the Olympian god and blacksmith Hephaestus fashioned automata, as well as Achilles’ shield. The Greek inventor and philosopher Heron of Alexandria wrote a treatise on automata; such devices were often cleverly set in motion by falling water, heat, or atmospheric pressure (15). According to Greek mythology, the inventor Daedalus crafted wings of wax to enable himself and his son Icarus to try to escape Crete, after Daedalus had built the Minotaur’s labyrinth there.

Icarus' wings were melted when he flew too close to the sun. Daedalus fashioned automata as well as prostheses; he is said to have built a copper machine, Talos, to guard Crete.

The Golem

The medieval legend of the Golem, a man-like figure formed of earth and called into animation by the invocation of the name of the Lord (i.e., a *schem*), was a late-medieval essay in alchemy. According to mystical tradition, the Golem was created by Rabbi Loew of seventeenth-century Prague to protect the Jews during periods of persecution. As the Hebrew word *Adam* ('from the earth') itself states, man is made of earth, as is the Golem. But the Golem legend emphasizes the singularity of the Divine in the ability to create life, as the Golem is mis-shapen where mankind is not. The idea has echoed in every subsequent portrayal of anthropomorphic robots, most notably in Mary Shelley's *Frankenstein* in the 19th century. The concept of the Golem has apparently been highly attractive to AI's founders as well: John Von Neumann, Norbert Wiener, and Marvin Minsky all asserted their direct descent from Rabbi Loew (16).

The Ars Magna

The Golem was a legend, albeit one with considerable and lasting psychic reality. During the Enlightenment, efforts to conduct mathematical operations with machines would appear, as would the idea of conducting such operations to systematically produce ideas. Curiously, the earliest effort to generate logical statements systematically using a machine was invented far earlier.

The *Ars Magna* (or *Great Art*) was invented in the Thirteenth century by the pre-Reconquista Catalan theologian Raymond Lull (1232-1315). It was a tool which could generate all combinations of a limited number of axiomatic principles, or concepts which were "true." These 'true' axioms were the patently desirable Catholic virtues or divine attributes- goodness, greatness, and eternity, etc. Lull's invention was a wheel made up of two or more

concentric circles. Each circle contained the fourteen accepted divine attributes. By rotating the circles, every potential different combination of factors- that is, one hundred and ninety-six twofold combinations- could be generated. Each of these elements could be combined with every other element to produce an exhaustive inventory of all true statements. The generation of combinations was syntactic rather than heuristic. All of the combinations, and not some selected or constrained result, were presented. The device avoided the eternal computational problem of a combinatorial explosion, simply because the very small number of inputs. Lull constructed similar tools for studying the seven deadly sins and other theological artifacts (17).

Because the Ars Magna presented different permutations rather than any novel knowledge, it was a pseudo-computational device. Thus it epitomized rather than transcended the Medieval dovetailing of theology and science. Notwithstanding this, it was the very first effort to think systematically using technology.

The Concept of Symbolic Languages and Thinking as Symbol Processing

The first person to envision symbolic computing, the ultimate sine qua non for AI, was Wilhelm Gottfried Leibniz (1646-1716), inventor of calculus (along with Newton) and of the Stepped Reckoner calculator. One of Leibniz' most intriguing theories is that of monads, atomic bits which express tiny aspects of given philosophical principles. This concept offers the idea of manipulating symbols systematically. Just like English philosopher Thomas Hobbes a hundred years earlier, who had declared much more briefly that 'all thinking is but ratiocination', Leibniz lacked practical digital computing and computer languages to demonstrate his obvious grasp of the concept.

In the absence of any objective economic need for a computer, the hints that both Leibniz and Hobbes proffered as to symbol processing remained as philosophy rather than computing. AI needed general-purpose digital computing, which was obviously far from existence before the Twentieth century.

Computing, Practical rather than Symbolic

The digital computing lineage, involving the manipulation of information through binary codes made up of rudimentary items such as punch cards, was intermittent. The practical lineage of computing, which did not intersect with Enlightenment philosophy, is just that- practical. Computing in the sense of processing large volumes of information mechanically was actually borne out of a need to weave cloth and to figure out rows of numbers for businesses and governments- that is, economic need, not intellectual curiosity. If we think of computers as tools to handle the raw manipulation of data itself, then we see that the history of computing originates much earlier than anyone might anticipate. Weaving patterns are the primeval form of programming. Paper or cloth markers provided the first mechanical means for processing (relatively) large volumes of information in the form of rows of thread. Punched cards started with the need for "easy storage of large amounts of information to be read not visually but mechanically"(18). Punched card usage was perfected by Joseph Marie Jacquard (1752-1834), who created an automatic version of the process for storing woven textile patterns on punched paper cards. Jacquard was further inspired by another French inventor, Gaspard de Prony, who found the idea of dividing large calculations amongst groups of less-educated workers in Adam Smith's division of labor (19).

The punched card form of storage was so successful that both Charles Babbage, designer of the first general-purpose computer, and Herman Hollerith, the American engineer who founded IBM, used abstract programs stored on lightweight cardboard pieces. On the cards used by Hollerith, punched and unpunched sections conveyed differing light signals (20).

We can trace a direct, if broken, line from Babbage to the first digital computer, in the mid-1940s. British inventor Charles Babbage (1791-1871) was the first person who tried to build a digital computer, that is one in which information is held in a binary state code (either on or off), rather than a decimal code, and which is therefore immensely flexible. He nearly succeeded.

His first computer, the Difference Engine, was a design for calculations to solve large polynomial equations. It was not built, for reasons which are best summarized as political and managerial, rather than technical. Working with Lady Ada Lovelace, who envisioned computer programs to run on computers, he designed a further machine that would sort cards according to binary signals. Babbage's second design for a computing machine, to be called the Analytical Engine, was also never built, again more for a want of management skills than because of technical impediments (21).

If the idea of digital computers continued to fascinate, so did the concept of automata. The art of automata had fallen latent in the European Middle Ages, and revived during the Renaissance. Scholar Roger Bacon (1214-1292) is said to have fashioned automata as well. Leonardo de Vinci thought of animals as complicated mechanical systems, that is, essentially automatons. The field of mechanical automata flourished in the Early Modern Era, as fantastic figures intended as parlor games for the aristocracy were much in demand (22). These were, like Babbage's work, luxuries and experiments. Despite Babbage's vision, there was no vital economic need for digital computers in the mid-19th century. They had been glimpsed, albeit repeatedly, before their time. It was not so for the analogue side of computing. Analogue computers measure quantities rather than abstractions, and are therefore well-suited for counting and mathematical activities, and measurements of continuous functions (for instance, as in the mercury thermometer), fit within this line of activity. In the Twentieth century, computers turned from aristocratic parlor games into working accounting devices indispensable for industry, as the digital line was finally taken up again with the ENIAC project during the Second World War (23).

Conclusion

The thematic effort of AI in its formation was "the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it." Having said this, we hasten to our story,

which will start with the state of the art of thinking about computing, intelligence and machinery in the Thirties and during the Second World War.

Chapter 2. The Ether of Ideas in the Thirties and the War Years

Introduction

“ Where there is a will, there is a way”.

So it is said, but it's more complex than that. The will for AI's existence was present in the early part of the 20th century. Rossum's Universal Robots, Karl Capek's famous fictional work, was not an isolated opus. Science fiction stories, featuring spaceships, talking robots, and voyages to the Moon, which were then becoming popular, were fanciful rather than real. The way did not yet exist, though.

AI as it would be realized in the statements of the Dartmouth Conference required working general-purpose digital computers; software programs; and an idea, however basic, of how human intelligence worked. The state of the art in the years leading up to the Second World War lacked these things. However, this decade did provide the antecedents which would become practical once the computer had been developed. A variety of disciplines studied intelligence in machinery and in the abstract, without ever discussing intelligence in people. These include Cybernetics, a theory of intelligence in machines with a good deal of abstraction so that it could be applied to humans as well; neurology; information theory; formal logic; automata (primitive robotics); and theoretical automata, or machines consisting of formal language or logical programs that could emulate life functions.

1. Before the Computer: Intelligence as a Railroad Routing Station

As we saw in the first chapter, thinking about the human mind, and about human intelligence embodied in machinery, is a

constant in human history. In the absence of raw materials, people will create stories about automata. Given scant raw materials, it seems, people will build automata. In the Thirties, practical automata were a popular hobby if not academic study (24). These electromechanical automata emulated biology rather than cognition. Science fiction typically predicted mechanized intelligence, in which highly anthropomorphic robots would carry out intelligent activities. This was presented in a 'black box' form, without any explanation of the basis, computational or otherwise, which would be used to implement it (25). There was no general-purpose machine that did 'intelligent' things. The shell of the idea of a computer existed, but there was no reality to substantiate it.

The idea of intelligence involving information processing or a program- that is, the computational metaphor for intelligence- did not and could not exist. It would not for years after the computer itself was extant and functional. Even as late as 1949, electrical engineer Edmund Berkeley does not venture far from the popular misconception when he suggests that the reader think of a mechanical brain as a sort of railroad station, in which information and sensory data are most processed by being routed to the correct location (26). It is not surprising that Berkeley uses this terminology- which discusses the routing of people and materiel- rather than the processing of information or goods. The latter would be an appropriate way to think metaphorically of computing, but metaphors of intelligence at the time did not support such a computational understanding. Intellectual historian Murray Eden discusses the weak metaphors used to describe intelligence:

“ In their introductory psychology text (1921, 1949), Robert Woodworth and Donald Marquis make no use of such concepts as message, information feedback, computation or control. A discussion of neurology is included but its relation to psychology does not go beyond the description given by Dunlap at least 35 years earlier. For these authors, the stimulus/response paradigm is preeminent... Woodworth and Marquis recognize the importance of organization but offer no means of investigating this property of the brain... We note again the analogic use of

telephone cables and explosives. I find no reference to the analogy of the computer. Furthermore in the chapter on perception, there is no bibliographic entry later than 1943. It is fair to conclude that the authors did not believe that perceptual research during the twenty-four years prior to this edition was of particular import to beginning undergraduates.” (27)

Eden explains that the authors lacked the terminology in terms of which to describe mental functioning apart from the transmission of neural impulses- in both 1921 and 1949.

2. Warring Camps in the Field of Psychology

Psychology concerning thinking was deficient during the entire first half of the 20th century, especially during the arid InterWar period. The intellectual state of the art was characterized by Freudianism, early Behaviorism and its compatriot scientific management, and progress in neurology. Cognitive psychology, the systematic study of thinking, simply did not exist. Both Freud and the school of academic psychology known as Introspectionism studied the mind, but it did so mostly as regards emotions. Freudian psychology, studied in psychoanalytic institutes rather than in universities, held that the mind was all affect, and the academic discipline of Introspectionism that the mind was hardly knowable (28). A little Freud is a good thing, but this was carried much too far in both academic and intellectual circles and in the culture of the intelligentsia. Whatever progress Freudian psychology may have made to the nascent state of the art in developmental and affect psychology before its advent, it unfortunately detracted from efforts to understand the nature of thinking as cogitation. Despite Freud’s arduous training in physiology and his practice as the most dogged of scientists, in this respect his work played into the hands of anti-technological and even anti-intellectual sentiments (29).

Behaviorism at its Most Radical

Behaviorism, which triumphed and was indeed ubiquitous in academic psychology for half a century, smothered the study of thinking. Behaviorism insisted that human beings’ behavior could

be predicted in a highly regular manner without reference to thinking (30). John Broadus Watson (1878-1958), the field's founder, declared that psychology was as cut-and-dry as wrapping a gift:

“...it is a purely objective experimental branch of natural science.” (31)

This school discarded the evidence of cogitation, consciousness, emotion, deliberative, the ‘internal life’ of human beings in favor of animals as research subjects, and fixed and predictable inventories of responses to given stimuli as the scientific evidence to which all psychology should be reduced. Harvard professor B.F. Skinner maintained immense sway over the research agenda (32). Much academic psychology consisted of quantifications of the reflex arc, or idealized process by which people respond to given stimulus. The experimenter tried to find clear stimuli, which would elicit, consistently and measurably, a particular behavior. Experimenters typically used Classical or passive conditioning, as in Pavlovian stimulus and response tests, which elicit salivation. However, Radical Behaviorism, Skinner's forte, relied on operant conditioning, in which the test subject is asked to do something. Operant conditioning was used increasingly as Skinner became more influential over the decades of mid-century. Because human and even much animal behavior necessarily incorporates far more complex information processing, this line of research was doomed to stall in its own limitations.

The two polar opposites of Behaviorism and psychoanalytic psychology, clenched together in eternal disdain, slowed work in cognitive and physiological psychology through at least 1950. The severe differences between these fields show no respite as we witness the intellectual world circa the late 1940s. Psychologist Howard Gardner's authoritative history of the development of cognitive science clarifies the chilling and limiting effects of the astringencies of Behaviorism:

“ So long as Behaviorism held sway- that is during the 1920s, 1930s, and 1940s- questions about the nature of human language, planning, problem solving, imagination and the like could only be approached stealthily and with difficulty if they were tolerated at all.” (33)

3. Formal Logic and the Universal Turing Machine

Cognitive psychology was essential to AI. Until it was pursued, even the basic premise that mental phenomena could be studied was not effective. Yet while it languished, formal logic moved forward. Formal Logic, especially productive during the Thirties, was an essential infrastructure for the development of computer languages. Its creators were philosophers, who worked independently from the electrical engineers. However, these strains would merge, very quickly, with the appearance of theories of intelligent automata and with computers. Logical positivism, a German and Austrian philosophical school, carried forward the advances of formal logic from the 19th century. It stated that real objects in the world were verifiable and could be clearly stated in a formal language. Formal logic attempted to describe items in the world in relation to each other. During the course of the 19th century, and into the early 20th century, the work of Frege, Boole and Bertrand Russell clarified formal symbolic references for belonging in a group, for qualities which shared common elements of two groups, the existence of objects, and some of their semantic (descriptive) properties. This level of formal logic would need to be elaborated greatly to allow for further description. Later, AI would incorporate formal logic in the Lisp language, invented by John McCarthy, as one of its major early intellectual advances.

The progress of the field of logic toward binary formal languages went as far as the Principia Mathematica, published by Bertrand Russell and Alfred Whitehead during 1907 through 1910. Russell and Whitehead developed a formal language for the expression of objects, but did not develop a rich semantics for their description. This achievement hung in the air with little progress until the 1930s. At that time, other logicians set a cap on

the further achievements of consistency in logic. But at the same time, logic also affirmed the possibility of computation.

David Hilbert, Schroeder, Russell and Whitehead, and the other mathematicians and logicians of the 'logical positivist' movement, had aimed to render formal logic, specifically first order predicate logic, as capable of expressing and indeed reiterating the objects in the real world (34). The caveat to that initiative was imposed by the Czech-German Kurt Godel (1906-1978), who trained in Vienna, moved to the USA in 1938 and settled to research at the Institute for Advanced Study. Godel's 1932 Habilitationsschrift (35) at the University of Vienna stated his most lasting contribution- the famous 'yes, but' which is known as the incompleteness theorem.

This work discovered that consistent logical systems may not always be able to verify themselves, nor to express every true sentence. These are Godel's Theorems, or the Second and the First Incompleteness Theorems of 1931, respectively (36). In addition to the finding that there are some- albeit mostly hypothetical- limits to formal logic's conquests, Godel's work also proved the completeness of first order predicate calculus. Several years later, the American logician Alonzo Church demonstrated that FOPC is undecidable (37). Like Godel's thesis, Church's theorem (1936) does not hamper the application of formal logic, specifically FOPC, to computer languages for use in workable technology (38).

The Universal Turing Machine

Prior to the era of workable digital computing, much of the essay to integrate logic into mathematics may be seen as castles in the air. Incredibly, these virtual artifacts became more and more realistic in their ultimate goals. The envisioning of automata, independently, by two thinkers at almost the same time, began as a far-off abstraction and ended up as guiding principles for the actual design of computers. Alan Turing and Emil L. Post both wrote about automata, in very similar terms, at almost the same time. "Finite Combinatory processes,

Formulation 1," (39) the paper by City College of New York professor Post, received less attention than that by Alan Turing (40). At the time, Turing was a visiting scholar at the Institute for Advanced Study, home of Einstein and Von Neumann, so this imbalance is not surprising. However, Post's paper is often considered an important prerequisite to the development of production systems, an important AI tool invented several decades later. With a nod to Post, we will consider Turing's paper, and Turing himself, in somewhat more detail. Alan Mathison Turing (1912-1954) made central contributions in early digital computing (the wartime Colossus computer), in the theory of automata and computers, and in definitions of the idea of AI.

Practically speaking, the impediments to logic pointed out by Godel and Church did not impede the development of computer languages, since the interesting thing is not proving that everything can indeed be computed, but figuring out the much narrower range of what is wanted. The 1936 work by Turing proposed the theoretical design of an automaton, that is, a machine that would process information in the form of marked squares of a paper tape without human intervention (41). We have previously discussed automata as an effort to reiterate the human body and its functions. The erudite discussions of automata during midcentury, which contributed a great deal to computing and AI, mostly concerned information processing.

This was the case with Turing's work as well. He designed the most elemental number-processing device he could. It could be made of paper, and proposed that this 'Turing Machine' could in principle carry out any computation that could be stated in algorithmic form. The machine is also called a Universal Turing Machine, as it is claimed that such a machine can simulate the workings of any other computer (42). This is a finite state machine, meaning that the number of states it may take is limited. The Turing Machine is a box with an input side, an output side, and an internal scanner. The operation of the machine is most schematic: the scanner can read one square of the paper tape, which is the form of input, at a time. The paper tape squares are marked with symbols from a limited code of 0s and 1s- in essence, a binary machine language. The machine can also

move the tape forward or backward, and can print new symbols or erase symbols. (More specifically, the machine instructions are “sets of ordered quintuples”). A mathematical function is computable, according to Turing, if the machine will be able to recognize, and stop at the function- that is, at $f(x)$. A mathematical function is said to be ‘Turing-compatible’, and hence suitable for computing not only on the UTM (as it is now called) but on other machines, if the Turing Machine will halt at the function.

The concept of the Universal Turing Machine raised the theoretical ceiling of what can be computed, but did not equal actually working with computers to develop programs. This is because instantiations of artificial intelligence will of necessity be much more limited. One’s automobile can indeed run at 110 mph; this does not mean that most owners ever test this limit. The importance of this paper to computing was immense. In the words of computing historian William Aspray:

“ The importance of the Universal Turing Machine to computer science becomes clear once it is recognized that it is a theoretical model of a digital stored-program computer.” (43)

4. The Curious Meta-Field of Cybernetics

In the absence of cognitive psychology, there was no way to talk about intelligent thought processes. However, there was a considerable need to develop intelligent engineered processes. In lieu of AI, cognitive science, or software engineering, Cybernetics provided the terminology during mid-century. As such it proved an essential bridge into AI and cognitive psychology.

Cybernetics, the pre-eminent way of talking about intelligence, was one of the oddest intellectual ducks of any time or place. Was this a theory of the nature of intelligence or the human psyche, like Freudian psychology ? Not really. Was it an empirically-based explanation of natural phenomena, like Darwinian evolution ? No, it was not that either. Instead, Cybernetics was a meta-science rather than a science strictly speaking. It described intelligent processes so generally that it did not rely on empirical research.

Cybernetics' most prominent founder, MIT mathematician professor Norbert Wiener, Wiener professed a lack of interest in "research topics which were amenable to the division of labor". Nice work if you can get it- and as a unique former prodigy, Wiener could. Very well, then, but he had to pay for his high-mindedness. Scientific research is almost always amenable to the division of labor. Scientific theory- however intelligent- which is not subject to being broken into the smaller tasks of designing experiments, carrying them out, conceiving models, clarifying one's place within a department, etc., is in danger of ending up as sterile philosophical hand-waving.

Wiener and his colleagues Arturo Rosenblueth and Julian Bigelow originally designed ballistics tables during the Second World War (44). They adopted the terminology involved in the mechanical systems of such weaponry- feedback loop, negative, positive feedback, and inhibitory feedback, into more generic usage. Wiener et al. proposed that this was universal terminology for intelligent processes; the authors' objective was to define the behavior of intelligent entities. A 'behavioristic' approach, as they called it, to assessing entities- without explicitly limiting the agent to human 'conscious' behavior- as the output of an entity in reaction to its environment, is defined in their earliest statement:

" The behavioristic approach consists in the examination of the output of the object and of the relations of this output to the input. By output is meant any change produced in the surroundings by the object. By input, conversely, is meant any event external to the object that modifies this object in any manner." (45)

The authors progressively refine the nature of the entity's behavior as a form of self-correction of output in response to the reaction of the environment. The study emphasized the methodological concerns of assessing behavior in terms of observable quantities such as input and output, rather than by (purported) internal design. These definitions of input and output were explicit, although highly abstract. By its high-mindedness, this scholarship thus provided a means of approaching complex systems, mental or engineered.

Purposeful or goal-oriented activity is emphatically so when it is non-teleological, or involves negative feedback. In a negative feedback loop, the entity responds to the environment's response to its (the entity's) behavior by in turn adjusting its program or behavior closely so as to better approximate the precisely designated goal. The authors specifically note that in this kind of behavior, the entity modifies its own behavior "in the course of the behavior." (46)

Finally, Rosenblueth et al. describe the behavior of an entity in terms of the inferred complexity of calculation which the entity performs concerning the path of the goal. A cat pursuing a mouse figures out the mouse's path. In so doing, it performs extrapolative or predictive behavior. The authors suggest that behavior which is based solely on sensory receptors is less likely to involve imputedly mental calculation and therefore less likely to be predictive. For instance, a bloodhound merely follows its nose, which senses keenly but predicts nothing (p21): "The integration of input and output necessary for the performance of a predicative reaction..." will only take place when the internal organization of the entity permits such coordination (ibid.). Extrapolative behavior implies that the entity in question refers to some internal cache of knowledge in order to 'figure out' how to respond to the environment.

The paper's ultimate concern is the internal mental and neurological organization of living beings. The authors acknowledge the limitations of applying generic models to vastly different categories of entities. That is, "the ultimate model of a cat is of course another cat" (ibid.). Once issuing this caveat, however, they proceed without hesitating. Thus the distinction between a complex process of prediction of the environment and simply reacting to the environment at every instant without any internal calculations:

"...leads to the singling out of the class of predictive behavior, a class particularly interesting since it suggests the possibility of systematizing increasingly more complex tests of the behavior of organisms. It emphasizes the concepts of purpose and of teleology, concepts which, although rather discredited at present,

are shown to be important. Finally, it reveals that a uniform behavioristic analysis is applicable to both machines and living organisms, regardless of the complexity of the behavior.” (47)

Rosenblueth, Wiener, and Bigelow's presentation of the concept of negative feedback was applied to mechanical devices and (implicitly) to human beings alike. The paper minimized- but did not ignore- the distinctions between human beings and machines that determine the moment at which a bomb will blow up an airplane. This instantly widened the complexity of tasks which machines were imputed to be capable of carrying out. The conclusions reached in “Behavior, Purpose, and Teleology” far surpass the technical minutiae of formulae for calculating aircraft trajectories, as well as greatly surpassing the state of the art in feedback control mechanisms.

Thus the trio, and soon a number of other colleagues, proposed the analysis of intelligence in human and other ‘complex’ entities. This paper formed much of the basis of Cybernetics as meta-theory, asserting all of the basic features which made it universal. The authors’ analytical framework sought out determination of the nature of feedback mechanisms, or learning, the degree to which such entities were capable of non-teleological analysis. Simply postulating the existence of intelligence as purposive was an explosive idea.

The potency of this theorem was evidenced in its universal impact during the PostWar period. This is particularly the case because Cybernetics offered to find these structures in human beings and in mechanical (engineered) objects. The discipline was presented as a universal one which could contribute something to practically all sorts of behavior, even across the apparently definitive line of organic versus inorganic or designed entities. In addition to its purported application to all possible objects of study, the field provided a general methodology and vocabulary. The universal target of these ideas allowed the concepts to spread throughout the sciences and social sciences in the PostWar era. The diffuse nature of the target meant that the field

had a perimeter and vast influence, but no core, and no real academic departmental home.

Surely one reason that the Cyberneticists took the neurological rather than psychological course to looking at intelligence was that they had been trained to approach the flesh and blood, or the frequencies, rather than the contents of thinking. The academic backgrounds of the Cybernetics people were almost always in fields other than psychology. Physics as applied to radar, and acoustical engineering, physics, electrical and mechanical engineering, and occasionally psychiatry or medicine were common. Many Cyberneticians- including Ross Ashby, Walter Rosenblueth, and Frank Rosenblatt- had been trained as physicians or psychiatrists (48). Thus, it is not surprising that there was almost no interest in cognitive psychology amongst this cohort. The emphasis upon distinct dedicated machines rather than software certainly fed the Cybernetic train of thought, which emphasized the machines' emulation of biology rather than the later AI approach to machines as 'thinking' entities. The biological approach persisted for the duration of the period through 1960, at which point AI and cognitive science were much more firmly underway.

The biological metaphors of Cybernetics also drew intellectual vigor away from the potential for thinking of software as a metaphor for thought. Because there were no computers by means of which to run computer software programs, there was no way of practically demonstrating theories of intelligence using computers, and less likelihood of understanding thinking as a protocol. Such theories could not be based on cognitive psychology experiments because the academic discipline did not exist.

However, Cybernetics was so openly concerned with the abstract concept of intelligence itself that it became helpful to practically everyone in the wide range of fields about the mind, intelligence, and computers in mid-century. Cybernetics never explicitly talked about the mind, but it allowed the discussion of intelligence. In this sense, its contribution as a bridge to the

explicit discussions of intelligence in cognitive psychology and AI a decade or so later was priceless.

5. The Origins of Information Theory

The Universal Turing Machine, as we saw earlier, was the concept of an abstract intelligent machine that could solve problems presented to it in 'software' paper tape written in a simple binary language. Its inventor, Alan Turing, was a logician and mathematician rather than a hardware designer. His 1936 'machine' did not closely address the nature of the coded machine's language, the scanner, its source of motive power, or other computer innards. There had been many advances in formal logic, but these were not effectively linked to running a computer. The gap between the idea of computer programs and computers for the practical purposes that World War Two demanded was still significant.

Wartime designers Atanasoff, Stibitz, Aiken and Zuse all sought to develop computing machines which were as universal as possible (49). Given these efforts at advances, computer hardware would have proceeded, regardless, at some point to the need for functional formal computer languages. This would require an intellectual insight into machine states and formal logic alike- specifically, how to write computer languages and run them on computer hardware.

Fortunately, by the time there were many computers in need of programming conventions, engineer Claude Shannon had already figured out how a very limited repertory of machine states could be combined to put into effect the basic Boolean logical operations. These in turn could be easily indicated in the machine languages of zeros and ones. The Cartesian or mind-body problem of computing- how the language and the problems it expressed would be effectuated by inanimate machines- was solved.

In 1938 Claude Shannon (1916-2001) was an MIT graduate student in electrical engineering, working under the direction of Professor Vannevar Bush on a project known as the Differential

Analyzer (50). This machine was an eminent example of analog computing. This sort of design, in which the design of the computer itself physically, rather than mathematically, models the relationship between the variables (51), has been all but entirely superseded by digital computing since then. However, it was the preferred way to proceed in Interwar period. The machine contained digital (that is, discrete-state) circuits which employed electromechanical relays. These circuits broke and required repair regularly. Shannon's insight into the isomorphism between information moving through circuits and conveyed through logical symbols arose from his repeated efforts to fix the circuits (52).

As expressed in a 1938 paper based on his master's thesis, Shannon's insight was that relay circuits corresponded exactly to the symbols of propositional calculus. (Specifically, the correspondences were: true /closed; false /open; and /serial; and or (inclusive) / parallel). What was needed was a way to program that translated machine states into logical rather than numerical cognates. Shannon provided this. The switches themselves had only the open and closed positions- an open switch would break the flow of current, and thus negate an add operation, while a closed switch would allow current to flow and thus was considered true.

As we saw in the beginning of this chapter, the early Twentieth century used the concept of railroad or telephone switching systems as means of information transmission, as well as a model of the function of the nervous system. This fit well with psychology in the first half of the century, which was more concerned with the transmission of neural information than with its processing as semantically meaningful data. The telephonic relay of data metaphor provided the grist for the development, by Shannon and MIT colleague Warren Weaver, of information theory. During the postwar decade, they further developed the concept of information carried through binary code, which labels and explains means of manipulating the electrical signals, which traverse telephonic and computational systems.

The theory is concerned with the eradication or correction for noise, or extraneous material, which invariably appears in channels of information being transmitted. Clarity of the message- any message- is of concern, whereas content is not (53). Information theory also provided crucial descriptive terminology in some areas, such as the invention of the concept of bit, or “binary digit”, the single binary choice, or the smallest possible unit of information which actually says anything. The latter central concept was conceived by Shannon’s colleague John Tukey of Bell Telephone Laboratories (54).

“ So a single bit resolves the uncertainty of a coin toss.”... “ There is an equivalent and more constructive way of getting at the number of information bits in data that we receive. It is also the minimum number of yes-no questions required on the average to arrive at the given data, knowing everything that we do know. Each bit of information is the answer to a single yes-no question...”(55)

By offering an algorithmic remediation to the noise characteristic of information transmission, Shannon made a massive contribution to the reduction of noise and hence errors in data transmitted. The introduction of checkbits and other forms of ‘coding’, to the message being sent, helps to assure that comparison between the message at origin and destination will reveal errors with a minimum of effort. The value of such error codes, developed by Shannon and others starting in the 1940s, to telephonic and digital networks was beyond calculation. (At that time signals were differentiated into analog and digital; effectively they are not today). Also beyond calculation for all transmission of voice and digital information, is the value of more parsimonious recording of information, through sampling (56).

Despite their earlier disclaimers, Shannon and Weaver tried to extend concept to semantics. It did not work. A probabilistic or statistical approach to information processing was of relatively little use to AI in the early years of the field (although later on it would be of much use to speech recognition and other forms of emulation of intelligence).

6. Conclusion

At present, we think of the mind as a sort of information processing machine. In contrast, the earlier approach was envisioned it- as a signal processing machine. With the latter metaphor employed as the standard terminology of information processing, emphasis would continue to be on computing hardware rather than on the contents of the transmission. Inadequate attention would continue to be granted to software; computing could not be standardized on any sort of hardware platform. Yet the increasing interest given to the semantics of information-processing in cognitive psychology and in automata studies meant that thought as information processing would become a focal point of interest. This would be a gradual process, over the several years following the creation of the digital computer at the end of the war.

The realization of intellectual projects that were made possible by the computer's very existence depended on a cohort of engineers, psychologists, mathematicians, and even, eventually people trained in computer science. This would require its own infrastructure. We will study this in the next chapter.

Chapter 3. The New World and the New Generation in the Forties

1. Introduction

The intellectual precedents for AI were starting to be formed- in formal logic and in ongoing increasingly elaborate discussions of intelligent automata- during the PreWar environment. However, this era did not develop more general-purpose digital computation. Unfortunately, that required a war.

The Second World War pulled technology forward, with no pause, in grand and terrible ways. At the start of the war, the digital computer had been envisioned but never built, and there was no compelling reason to sink the requisite enormous energies into the project. However, the project was developed for the

calculation of missile trajectories at the Ballistics Research Laboratories at Aberdeen, Maryland (57). Once the mad rush it required had been put underway, it did not stop- even, ironically, once the German and Japanese forces had surrendered in May and August 1945 respectively.

The ongoing segue into the Cold War created a new cohort of tens of thousands of people who developed the new technologies and sciences which were associated with this war- broadly speaking- as well as its intellectual and logistical infrastructure. Among this enormous cohort were the future creators of AI, as well as their larger circle of intellectual compatriots.

2. After the War: The World Created Anew

The raw shocked air of the late 1940s was punctuated with new particles. The Second World War had torn vast swathes of the developed world to rubble. Practically everything was remade in the years following it. It is not surprising that the world of the late 1940s seems to have been little but novelty. It was seen in new materials and goods such as nylon and margarine, in innumerable 'temporary' buildings that stayed put for decades, in vast new suburban single-family housing developments, in the growth by leaps and bounds of the Sunbelt states, in the ongoing militarization of society, in the G.I. Bill and the associated rush to college among millions of former enlisted men. Western Europe, practically demolished, was no longer pre-eminent in economic and military power. The Soviet Union had taken military and political command of Eastern Europe. The United States was one of two world powers now, and the sole one which had not been invaded. The USA owed this novel power to its military might and victory in the war. It would maintain this position only by virtue of continued success in the Cold War. Plain-spoken and pragmatic as ever, Vannevar Bush stated, "If we are really well armed, the Reds will not force a world war on us." (58)

The phrase "really well-armed" could not have been more broadly construed in practice. The implementation of this concept

practically redefined the universities, research institutions, social and hard sciences, as well as the more obvious armed forces.

3. The Cold War and the Pursuit of Technology Research

“ The national effect of this mammoth effort has been enormous. A vaccine for German measles, the space program, freeze-dried coffee, solid-state electronics, color TV, lasers, LSD, the high-speed computer, automation schemes, the H-bomb, new breakfast cereals, Xerography and war contingency plans are just a few of the spectacular explosion of end products that have resulted from post-World War II R&D.”(59)

At least as much energy, if less blood, was spent fighting the Cold War as was spent fighting the Nazis and their co-combatants. This may seem odd because of the Soviets’ less malevolent aspirations- outside their own borders at least (60). Stalin claimed that his grasp was limited to “Socialism in one country”. However, he was exceedingly mendacious: the USSR was still adequately intimidating to keep up American military readiness. The Soviet regime endured longer than the Third Reich, by a factor of roughly six, and was internally and externally mendacious the entire time. Moreover, since the contenders were not continuously engaged in fighting a ‘hot’ war to the death, it was easier for both participants to wage a ‘cold’ one over a much longer period of time. Like it or not, this Cold War produced the phenomenal bounty of this Cold War. This is true of AI as well, although artificial intelligence does appear to be the most attractive child among the brood.

The Cold War could not be waged without dramatic redefinitions of the United States’ government. The United States’ military elite began planning its redeployment toward the ongoing enmity against the Soviets even before World War Two had formally ceased. At least one year before the Soviet Union began to swallow whole formerly independent Eastern European nations, the post-war mobilization began. Sudden bursts of research funding by the defense establishment have been attributed to provocations from the Soviets in the form of nuclear testing and

Sputnik. Its establishment as a whole was as much preemptive and strategic as simply reactive. The Soviets did not commence testing of the hydrogen bomb until the end of the 1940s (61). Yet Josef Stalin's naked aggression and utter amorality provided good excuses for preemptive militarization.

4. Science, the Endless Frontier, and the Inauguration of National Science Policies

“ We can no longer count on ravaged Europe as a source of fundamental knowledge. In the past we have devoted much of our best efforts to the application of such knowledge which has been discovered abroad. In the future we must pay increased attention to discovering this knowledge for ourselves...” (62)

In 1944, President Franklin D. Roosevelt requested that Vannevar Bush, Director of the Office of Scientific Research and Development and the U.S.' leading science policy figure, draw up a plan for the next generation of scientific expenditure. Bush responded with the famed white paper, “Science, the Endless Frontier” (63). He proposed a federally-funded group of institutions which continued along the lines of the OSRD, addressing various research needs, including medicine, defense, and the natural sciences. At the time federal research initiatives were quite limited, and were mostly confined to the area of public health (64). Federal institutions, located and administered without regard to geography, doling out grants to promising individuals and institutions as Bush proposed, were anathema to the prevailing philosophy of governance, even immediately following the Second World War (65).

Following massive controversy, the Federal government succeeded in initiating a program of national research programs, both separate from and under the control of the Armed Forces. Each branch wound up with its own agencies. The Vinson Bill of 1946 established the Office of Naval Research (66). The ONR subsidized dozens of research projects, some with tenuous relation to warfare itself. For an interval in the late 1940s, the Navy was indeed exceptionally prolific in this regard. As G. Pascal

Zachary puts it in his biography of Vannevar Bush, “The Navy was on its way to becoming the what Newsweek called ‘the Santa Claus of basic physical science.’” (1999, p329). Among the items on the list of fields under consideration in 1948 were “geophysics, astronomy, math, chemistry, underseas warfare, fluid mechanics, psychophysiology, biochemistry, human ecology, physiology, microbiology and psychology.”(67)

The debate over a national foundation to encourage science and technology continued, with President Harry Truman once vetoing the establishment of the National Science Foundation on the grounds that it was elitist (68). He was utterly correct, but the new United States- the Cold War superpower- relied on technical and scientific elites, not on populist egalitarianism. Elitist or not, the NSF was finally established in 1951. Controversial or not, many ongoing military-derived research projects rolled right along. The ONR funded many projects. The most ubiquitous was the JSEP (Joint Services Electronics Program), a catch-all fund for miscellaneous projects of interest to the armed services (69).

The Air Force established its own office, the AFOSR (Air Force Office of Scientific Research) to subsidize third-party research. The Air Force also had its own think tank, the richest and most elite of the numerous think tanks funded after the War. The Rand Corporation was set up as a subcontractor for Air Force projects, employing engineers recently demobilized from the OSRD. At first, Rand was Project RAND, a department of Douglas Aircraft Company, with an initial \$10 million contract. This project was followed by independence, incorporation and more contracts, many of them civilian (70). Rand would prove essential to the birth of AI, as we will see in Chapter Eight.

5. The Brave New World of Basic Research

Popular opinion, and even the opinion of those inside federal government and the armed forces, was wary of the idea of basic research itself. Early on, the NSF had defined basic research as “the exploration of the unknown”, or the pursuit of knowledge for its own sake. This is ironic, since the Cold War was better

characterized as an ongoing pressure cooker in which every 'cooked' idea was served at the military dinner table. More cynical definitions of basic research abounded: Charles E. Wilson, President Eisenhower's First Secretary of Defense, called basic research:

“...what you do when ...you don't know what you're doing.”(71)

Yet the military research budget flagged only slightly even during the late 1940s. By the time the Korean War began in 1950, the budget was back at its full WWII level. It promptly doubled to \$1.3 billion dollars, and did not diminish for decades (72). The contributions of the scientific community to the U.S. success in the Second World War thus segued into an unbroken flow of military spending devoted to science. The average veteran returned to work, and to college, after 1945. But the scientific community, particularly the elite, was only partially demobilized; the scientific capacity of the OSRD was simply channeled into a multiplicity of different institutions.

MIT had been the most significant university contractor for the military during the war, as well as the university most closely affiliated with leaders of science policy (73). Vannevar Bush, for instance, had been an electrical engineering professor and then dean of engineering at MIT before he left for Washington (74). The Institute maintained this status for the decade afterwards, as the site of Whirlwind and SAGE, the two mammoth computing projects of those years. Project Whirlwind, a military flight simulator endeavor started in 1944, and marked the beginning of general-purpose rapid computing, leading to the invention of magnetic core memory (75). Whirlwind led to Project Charles, which led to a permanent concern, known as Lincoln Laboratories, in 1951. The World-War II-era Radiation Laboratory, or 'Rad Lab', an important weaponry facility under the OSRD (76), was renamed (but not re-staffed), the Research Laboratory for Electronics, or RLE. (In Chapter 11, we will discuss MIT's role as one of the leading institutions in AI's early development).

The largesse that paid for the Cold War research meant that there was more to spare for the “hard”, that is, quantitative, social sciences. These had not existed at all at MIT prior to the war, and the linguistics and psychology departments were set up as auxiliaries to established areas of study (77). The psychology department was established by J.C.R. Licklider, recruited from Harvard’s psychology department to work on Project Charles (78). The psychology department was initially established as a theoretical afterthought to MIT’s contracts on making military machinery easier for people to use. The department started out as a branch of electrical engineering. Regardless of the subordinate pecking order, the new areas of study succeeded in attracting extremely sharp minds (79).

The militarization of academia continued throughout the next thirty years. The Cold Warriors were often consumed with their research with an intensity of sentiment that startled younger people. This is famously true of John Von Neumann and Edward Teller, either of whom may have been the model for the fictional 'Dr. Strangelove', but the sentiment was not limited to Eastern European expatriates. In 1972, Patrick Winston, the newly appointed director of MIT’s AI Laboratory, asked various MIT lab directors for advice on how to create a great research institute. Winston reports that Jay Forrester, the father of Whirlwind, was the only one who had thought seriously about the question. But Forrester’s motivation turned out to have been political as well as technological:

“ I...said something like, ‘Well I suppose in any case it wasn't hard to have a great laboratory because everybody must be highly motivated by the Whirlwind computer and the thought of putting together the first one.’ And Jay looked at me like I was the king of the fools and said (laugh), ‘The laboratory wasn't about the computer at all; it was about the protection of the United States against air attack from the Soviet Union.’ That remark always stuck with me. The lesson I learned was the way to have a great place is to have a passionate mission.” (80)

The avalanche of new people, new funding and contract opportunities, and new technologies and science with which to

contend helped to spur a dramatic intellectual leap in the sciences of the mind. In the decade after the close of the Second World War, tremendous intellectual ferment gave rise to the digital computer, stored computer programs and early computer languages and programming aids, renewed study of cognition in psychology, far more sophisticated neurology, and finally to Artificial Intelligence.

At the same time, the most sophisticated echelon of the military leadership began to encourage scientific work which was revolutionary in implications. Greatly enlarged engineering schools, and in certain cases, entirely new departments and schools within universities were established, for the next two decades. This includes CMU's Graduate School of Industrial Administration (1951) and the Electrical Engineering department's Computer Science division at Stanford (1964).

If the existing departments and schools were, for some reason, not appropriate auspices for the carrying out of computing and engineering research- then different institutions were set up. This includes countless new departments of Operations Research and Electrical Engineering and Computer Science. RAND in Santa Monica, the Stanford Research Institute a few miles north of Stanford, Lincoln Laboratories at MIT, research firm Bolt, Beranek, and Newman in Cambridge, NASA Ames Research Center southeast of Stanford, and similar government contracting research institutions, either as university-affiliated private corporations or as parts of universities, were established as a virtual 'shadow university'. Scientific think tanks and weapons facilities were built in mid-century on tracts of flat bare land, in locations of varying remoteness. These include IBM Almaden south of San Jose, Lawrence Livermore Laboratories forty miles southeast of San Francisco, Oak Ridge in Tennessee, and Sandia and Los Alamos in New Mexico.

The disruption of the established ivy-covered buildings reshaped the United States geographically as well as intellectually. The unabated and decades-long defense spending spree, the endless hiring of professors in every field even

proximate to computing, the new think tanks and fields of computing and electrical engineering, all made the Postwar scientific and technical environment a golden one for its denizens. The opportunities afforded the generation in general, and the inner circle and outer periphery in specific, truly did not exist prior to 1945.

We will sketch out the inner circle of the founders of AI, and then the dozen or so people who would work closely with them. Beyond this, there was a much larger cohort who did not work together as a team by any means. However, collectively the hundreds and thousands of people who entered science would be drawn into all of the many fields associated with computing can be referred to as a generation. They were not a closely knit group or work team, but a cohort born and educated closely enough in time that they had similar, and favored, paths through life.

5. The Inner Circle of the AI Leaders

The men who would lead AI during at least its first two decades- John McCarthy, Marvin Lee Minsky, Allen Newell, and Herbert Alexander Simon- were uniquely favored by history. By accident of birth, they had dodged the bullets presented to many of those close to them in age. They were too young to have served in the Second World War, or to have had their personal histories and careers marred by the Depression. They were scientifically educated Americans at the start of an astonishing surge in technological innovation, crowned by the birth and continuing growth of computation. Surrounded by inventive contemporaries, and with their work supported by the lavish funding of the United States military, they enjoyed ideal circumstances.

Finally, like the larger circle we will discuss momentarily, the founders of AI were positioned as young men at the very start of the Postwar demographic bulge, or 'Baby Boom'. This meant that they would repeatedly find themselves in a senior or managerial capacity as that enormous generation was educated and began working.

John McCarthy (1927-2011) invented both the LISP computer language and timesharing (the programming convention that permits many people and processes to simultaneously use the same computer), along with numerous AI formalisms and the term “Artificial Intelligence” itself. His heritage was stridently politically leftist and economically modest. His father, a longshoreman’s union organizer with little formal education, moved to Boston from Ireland. His mother was a Lithuanian Jew, a college graduate, a journalist and a suffragette (81). The family was atheist, Marxist, and highly literate. McCarthy never doubted his own identity; as he told biographer Rodney Hiltz, “It was obvious to me that from the age of eight or ten I was going to be a scientist” (82). McCarthy was apparently always motivated by questions of his own, and recalcitrant to doing what he did not wish to. At times he carried this to extremes. He began his studies in mathematics at the California Institute of Technology in 1944, but was expelled because he did not attend physical education class. Once drafted, he sat out the end of World War Two as a supply clerk. He returned to Caltech and graduated in 1948, and then proceeded to Princeton for his Math Ph.D. (1951).

Marvin Minsky (1927-) has been a scientific polymath and tinkerer all of his life. Minsky was affluent; his father was a professor of ophthalmology at a university in New York City. He attended the Bronx High School of Science, and then finished high school at Phillips Exeter to help assure admission to Harvard. In college he nominally took a major in physics, but this was secondary to eclectic work in every scientific department. He followed this with a mathematics doctorate at Princeton (1954), and then, with the support of Claude Shannon, Warren McCulloch, and Norbert Wiener, Minsky was invited to be a Harvard Junior Fellow (83).

Allen Newell (1927-1992) grew up in San Francisco, the son of a professor of radiology at the Stanford Medical School (84). A graduate of San Francisco’s elite Lowell High School, he attended Stanford University, studying with the emigre mathematician George Polya (of whom more later) and graduating with a major

in physics in 1949. After a year sheep farming in Montana, Newell went to Princeton to study mathematics, but found graduate school “boring” and left after one year to take a research post at Rand. There he met Clifford Shaw, who later programmed Rand’s computer. Newell was also drawn into ‘man-machine integration’ research, concerning improvements to computer usability. This, in turn, led to his collaboration with Herbert Simon. In 1954, Newell moved to Pittsburgh to work with Simon. His research was submitted as adequate for a doctoral degree at Carnegie Tech’s Graduate School of Industrial Administration (GSIA). Once committed to the idea of artificial intelligence, he was utterly single-minded. Newell’s Grail Quest was defining the canonical atom of intellectual activity; this project and its numerous expanded nodes occupied him for the remainder of his life.

Herbert Simon (1916-2001) was the most patrician of the four, as well as being the one with the broadest interdisciplinary reach. One of the few human beings world-famous across academic disciplines, Simon was awarded the 1978 Nobel Prize in Economics. He took a highly circumlocutions route to get to Artificial Intelligence, or as he called it, ‘Complex Information Processing’. The son of an engineer who designed control gear, Simon was increasingly aware throughout his life of a familial destiny to broadly apply engineering to the human as well as the material sciences (85). But initially, he did not even study engineering or science, but actually majored in political science. When Simon was an undergraduate at the University of Chicago in the 1930s, digital computing did not exist. His first published works concerned techniques of Municipal Administration; he earned his doctoral degree in Political Science at the University of Chicago as well (1940) (86).

The circle of men shared various traits, increasingly so as their lives converged upon a common object. McCarthy’s origins were proudly working-class, but this did not impede his education. The others were affluent and scientific in heritage. The fathers of Minsky and Newell were university professors in science, while Simon’s father was an engineer and patent attorney. These

generally bourgeois socioeconomic backgrounds contributed to the young men's freedom to go to college and graduate school.

The courses of study which were readily available and encouraged was in itself a fortunate historical novelty. This generation was perhaps the first in U.S. history raised in relative affluence, made useful by the absence of elite disdain for practical technology. None of these men was from the East Coast wealth that has traditionally held political power and devoted its energies to government or law. The bourgeois, but neither elite nor effete, backgrounds and technical educations of the four afforded historically unique opportunities to become scientists. None of these people was a 'gentleman scientist'. They could pursue practicality in science without embarrassment.

In addition to good fortunes in background, the founders of AI were lucky in the timing of their birth- they could not have done better had they planned it. They were young enough to benefit from the first flood of postwar monies to higher education in technology and science, especially computing, and to escape being caught up in the political snares of the 1930s or the War itself. Three of the four were born within five years of each other. They were close to draft age at the very end of World War II, but not old enough to face any real danger. Minsky volunteered for the Navy in 1945, but the war ended before anything happened to him (87). (Curiously, AI inner-circle members Oliver Selfridge and Douglas Engelbart also happened to join the Navy within a few days of V-J Day). Simon, somewhat older, has a different profile: he arrived at computing as a tenured professor in his forties; he was declared exempt from military service because of colorblindness (MOML).

If these four were companions, colleagues and often friends, they were not 'fellow travelers', referring to participants in the American 'Old Left', that is, the Communist Party before the Second World War. Only Herbert Simon actually could have done much as a "fellow traveler". Instead, he became a welfare statist, and quiet but a lifelong advocate of civil liberties (88). Otherwise,

the evidence suggests a pronounced lack of interest in politics on the part of McCarthy (since the early 1970s), Minsky, and Newell.

But even those who are agnostic toward the political conflicts of their times are influenced by those times: this is especially true of this group. The Cold War was fought by weapons designers (among others), and the practitioners of AI, however personally intellectually independent, were among the civilian beneficiaries of the War.

6. Common Circumstances and the Larger Circle

These four men composed the nucleus of a wider circle of intellectual companions in the field of AI and cognitive science during the third and final quarters of the Twentieth century. This larger group includes perhaps ten other scientists and administrators who were closely affiliated with AI. This group includes Oliver Selfridge, who may very nearly be considered one of five, rather than the typically understood four, founders of AI; cognitive psychologist George Miller; linguist Noam Chomsky; and computer scientists Herbert Gelernter, Ray Solomonoff, George Forsythe and Nathaniel Rochester. Psychologist J.C.R. Licklider, an early and forceful administrator of government programs for AI, was originally part of the SAGE project and was firmly within the intellectual cohort of the AI scientists.

More peripheral to AI as such, but similarly brought into technology through work at think tanks (in this case SRI) was Douglas Engelbart, inventor of the computer mouse. The same is true of Clifford Shaw, whose working-class origins precluded elite formal education. This meant that he learned computer science as a working engineer at the Rand Corporation, and in the process worked closely with Newell and Simon on the earliest problem-solving AI programs. Herbert Simon, ever gracious, always accorded Shaw immense credit in publications through the 1950s. Shaw is central to AI's foundation; his participation only waned in the 1960s.

In addition to the ambient environment of a nation at peace, this group was given many reasons to devote itself to the technical issues of computing and its applications. These opportunities were offered in the form of direct mentorship or general intellectual guidance, employment at universities and think tanks, ready publication and monies facilitated by eager grants agencies and universities, and renown. Combined with the great intellectual aptitude of the AI founders and their cohort, these circumstances have provided grist for the intellectual mill for entire careers.

If material circumstances were made comfortable for this group, their intellectual heritages provided encouragement without debt. Directly and indirectly, each of the founders of AI was helped by Cybernetics and other contemporary intellectual developments. This encouragement was variously distant and impersonal, or personal and sustained. John McCarthy spoke with Von Neumann after the latter's talk at the Hixon Conference, and Von Neumann responded with encouragement when McCarthy wrote to him after the Hixon Seminar with his early ideas on automata. Allen Newell attended a talk by Claude Shannon while Newell was at graduate school at Princeton in 1949. Newell knew about Shannon's practicable method of transforming data into a binary logic-based language. But despite his interest Newell was reluctant to introduce himself to Shannon: "I was just a lowly graduate student"(89). Marvin Minsky was personally acquainted with many of the chief participants in the Cybernetics movement, as well as with their nemesis, B.F. Skinner. Minsky attended Licklider's seminar at Harvard in the 1940s, studied with Von Neumann at Princeton and elicited Von Neumann's approval on his doctoral dissertation, and won obtained the support of Warren McCulloch and Claude Shannon when the time came to apply for a Harvard Junior Fellowship.

Herbert Simon, a political science professor at the University of Illinois in Chicago during the 1940s, was geographically distant from Cambridge, at the time the undisputed center of attention. But he was well aware of the movement; he knew about Von Neumann and Morgenstern's work on game theory, and about

logical formalisms. (Rudolf Carnap, one of the pre-eminent logical positivists of the Vienna Circle, had fled Europe for Chicago) (90). In 1950 Simon wrote a paper applying the basic concerns of control theory to social and political organizations (91). This was two years before he began trying to emulate thought processes using computer programs, and his consideration of Cybernetics was obviously guarded. As a political scientist and economist, with some familiarity with symbolic logic, he immediately saw thought processes as more closely linked to symbol than to neurons. Simon's arrival at Rand gave him a crash course in the state of the art. As they got underway with the earliest cognitive simulation by computer in 1952, Newell and Simon knew of Von Neumann's insistent pondering of the conceptual similarities between computer and brain or neurological intelligence, and also about Oliver Selfridge, who was addressing the problem of pattern recognition at Lincoln Laboratories (92).

J.C.R. Licklider (1915-1995), a few years older than Newell, Minsky and McCarthy, was more directly involved with every aspect of the Cold War computing and Cybernetics experimentation. He began as a lecturer at Harvard, then became a SAGE engineer, and joined MIT's psychology department as the latter was set up with funding in part from SAGE's psychological corollary projects (93). He was also a participant at the Macy conferences, in a weekly meeting of several dozen Cambridge intellectuals held by Norbert Wiener, as well as in a similar faculty group at MIT (94).

The heritage of the Alinner circle was clearly marked by the cybernetic idea of intelligence in artifacts, but apparently not with a dogmatism that constrained later creativity. This was not merely a fortunate generation- it was provided perhaps more opportunities than any generation in human history. As we will see throughout our story, the succession of unique and brand-new opportunities meant that McCarthy, Minsky, Newell and Simon repeatedly were present at the creation of new fields, using new research tools.

Chapter 4. The Social Environment of the Study of Intelligence in the Postwar Years

1. Introduction

As we saw in the last three chapters, neither the computer nor a language to discuss intelligence existed in the years before the Second World War. Yet following the war, the computer and the creation of intelligent tools to improve it were suddenly immensely important. The social environment in which to improve such ideas continued to grow in intellectual intensity and population. AI was thus gestated in a highly energetic, optimistic atmosphere.

2. Social Synergies among the Scientific Elite

In the scientific fields with which we are concerned, the postwar intellectual climate was strikingly vigorous. News of the ENIAC- EDVAC projects spread quickly. The Moore School, site of the ENIAC, received visitors from the top of electrical engineering fields, from physics, and from mathematics. Domestic visitors arrived from many venues- the IAS, the NDRC Applied Mathematics Panel (the Wiener-Rosenblueth-Bigelow project discussed in Chapter Two), MIT's Servomechanism laboratory, RCA's research lab in Princeton, and Los Alamos Theoretical Physics Division (95).

The occasional meetings societies of the time perpetuated Cybernetic research and theory. The Teleological Society, formed at the end of 1944 in Cambridge by Howard Aiken, Norbert Wiener and Von Neumann, was intended to support the discussion of "communication engineering, the engineering of computing machines, the engineering of control devices, the mathematics of time series in statistics, and the communication and control aspects of the nervous system" (96). Five years later the Ratio Club, the British counterpart to the Teleological Society, was formed for comparable purposes (97). The Ratio Club ejected anyone who reached the rank of professor, putatively to try to prevent ossification. Its membership, led by Grey Walter, neurosurgeon John Bates, Alan Turing, electronic engineer Albert M. Uttley, and physicist Donald MacKay, was somewhat more

directly involved with automata- as actual machines rather than as theory, as Von Neumann and the Wiener-McCulloch-Pitts team were. The Club remained highly concerned with the brain as a model for automata and computing machinery. Only Uttley and Turing actually did work on computers (98).

3. The New Generation in the Early 1950s

The learned societies we just discussed were limited to members of the scientific establishment. In addition to these erudite high science meetings, there were less exclusive forums. Forums open to students included J.C.R. Licklider's seminars on Cybernetics at Harvard in the late 1940s (99), and the open public lectures at the Hixon Seminar at CalTech in 1948 (100). The Macy Conferences, which were open to very few people, were influential through those people (101). The same is true of the Hixon Seminar: John McCarthy, a CalTech undergraduate at the time, began thinking about Cybernetics after attending the Hixon lectures and briefly speaking with Von Neumann. The Hixon Seminar was even more directly consequential in inspiring cognitive psychology (102).

Other intellectual purviews, such as the defense-related Project Charles, Project Lincoln and Project Cape Cod, all at MIT around 1950, and the SRL at Rand and March AFB in the early 1950s, drew in larger numbers of people. The relative permanence of some Cold War projects was valuable: Lincoln Labs, Charles Stark Draper Labs, the long-lived Whirlwind project, the RLE and the Radiation Laboratory, Aiken's shop at Harvard, and Columbia University's Watson labs (and the associated Service Bureau), hired hundreds of people, some of them for decades. These job placements and meetings were discriminating and involved various levels of secrecy. The formation of the Association for Computing Machinery, for instance, took place in an open meeting at Columbia University, but one of the society's first computer conferences sessions were held at Oak Ridge, which was a closed city at the time (103).

The legions of people working on hardware in the 1950s were the prerequisite for the progress into reliable hardware, and hence work on reliable software. Apart from participation in these meetings, progress in computing continued. Many self-taught people and those who were somewhat younger, were simply not included or not even aware of the existence of these groups. The appearance of full-time jobs in engineering fields swept in hobbyists and home tinkerers, and turned their interests into careers. At this point, these careers were in hardware; applications software and systems software only really became viable in about 1960s.

Computer training itself became increasingly widely available. This included training under the auspices of the Armed Forces, which maintained a draft until 1975, other government agencies such as the Department of Energy and the Census a widening number of corporations, and popular books and magazines available to a growing audience.

The founders of AI were both generally in the Cold War technological environment, as well as specifically and personally in the circles of this environment. The specific intellectual involvements of Minsky, McCarthy, Newell and Simon were directly with Von Neumann, Nathaniel Rochester of IBM, and Bell Labs in the case of the first two, and with the Rand Corporation in the case of Newell and Simon. We will present the intense and ongoing discussion of automata, the concept of intelligence, and how it can be embodied, in the next several chapters.

Chapter 5. Von Neumann, Turing, and Abstract Automata

1. Introduction

Prior to functioning digital computers which could be programmed using software programs, automata provided a means to discuss intelligence. In Chapter Two we saw that physical functions were emulated in physical automata- in robotics, generally in non-anthropomorphic robots. Both literary depictions and actual robotic automata were popular during the

second quarter of the 20th century in particular. In this chapter we will consider automata which were developed to undertake more ineluctable and abstract tasks, that is computer programs before programming languages of any kind existed. These are essential missing links in the origins of AI, since they provided a way to 'build' intelligence prior to the existence of software programming languages.

The acquisition of the habit of referring to problem solving as something that could be illustrated by a software program had to be learned, over years of usage of early digital computers by AI's founders. The latter group, in turn, spent the 1940s and early 1950s in the intellectual environment of the era, with automata and Turing machines and logical formalisms as the preferred topics of discussing intelligence. The unique contributions to automata of neurologists Warren McCulloch and Walter Pitts, and mathematician John Von Neumann and logician Alan Turing are crucial in this respect. McCulloch and Pitts proffered a way in which intelligence could emulate the transmission of impulses by nerves. Their specific references to neurology were later disproved, but the concept of the correspondence of formal logic statements with computing clearly presaged software languages (104). Von Neumann, in particular, inspired the younger generation, by encouraging Minsky and McCarthy in some of their earliest work on the concept of intelligence embodied in an artifact.

Von Neumann and Turing, who we will discuss momentarily, worked in a highly theoretic manner. This is especially true of Von Neumann. Turing himself engaged in code-breaking and cryptography during the war. However, discussion of intelligence in extremely broad terms included the gamut from a cellular automaton or universal Turing machine that 'existed' more than a decade before it could be built, to early robotic automata. The extremely creative nature of these enterprises was encouraging. Yet the fact that they lacked an overall vocabulary or standard means by which to discuss how to build intelligence was problematic.

2. Von Neumann's Conceptualization of the Intelligent Automaton

An intellectual high note of the early 1950s was struck by John Von Neumann, in work which crowned his unparalleled career. Von Neumann was a first among equals in both digital computing and automata. His most important contribution, for our purposes, was the elucidation of the automaton. Von Neumann had taken up the idea of automata in 1936 after seeing Alan Turing's presentation of the idea of a logical automaton. As we saw in Chapter Two, this was a binary coded logical systems that could be programmed to carry out any operations including reproducing itself. Turing suggested an automaton based on the manipulation of binary digits, which could perform Boolean operations and hence take in instructions, including instructions to rebuild themselves (105). Today, we would call these binary digits 'bits', although Von Neumann did not use the term.

Von Neumann took up this concept, as is evident from his contribution to the ENIAC. He had been drawn, like a moth to the flame, into this design. In this capacity, he had considered the physical stuff of vacuum tubes and their earlier form of mercury delay lines. Yet he was flippant and dismissive of the weighty engineering challenges of early digital computing. Even in his first, undiplomatic Draft Report on the ENIAC, Von Neumann replaced references to Presper Eckert's logical schematics with allusions to neurology (106). For instance, a logic circuit was referred to as if it was a nerve firing. Von Neumann's interest in the implementation of logic and hence of information processing through what were at the time only theoretical automata, rather than real machines, continued and deepened. Like most of Von Neumann's work, this was solely conceptual. He focused his gaze on the design of a machine which could, given adequate operating instructions, run by itself. But he was not interested in the physical stuff of the automaton or of the digital computer, of its memory or logic units, or of the arduous actualities of its overheating problems or materials to build it with.

Von Neumann became increasingly interested in automata which emulated some of the functional features of living things, and less and less concerned with symbolic logic. He thought more

biologically, in a way that clearly prefigured A-Life and robotics. He had originally conceived of automata as logical binary devices, as Turing himself had suggested. This conceptual artifact, the early Universal Turing Machine (UTM), is ancestral to modern software architecture. But Von Neumann's new conception of the automaton differed from a UTM and ended up inspiring A-Life, much further down the road. The 1948 Hixon Seminar in neurology appears to have led Von Neumann to believe that the smaller informational units making up such an automaton should be cellular rather than logical (107). Again, this question was posed at a theoretical level. Von Neumann pondered the functions which such an entity must be designed to carry out. He considered the comparative way that a biological entity carries out these functions, and how they could find correlates in artifactual entities. This led him away from logic, and toward mathematics involving calculus (non-linear partial differential equations), to describe the automata. It also led him toward considering the biological nature of cellular communication, that is the biochemistry of the central nervous system (108).

Von Neumann's essays, presented at the Hixon Seminar, evidence that he was much more impressed with the brain than with the computer. The work ends up being a paen to the brain. Like the Yale Silliman lectures, which Von Neumann prepared but could not deliver (1958), several years later, "The General and Logical theory of Automata" paper takes the body as the template and the computer as the model copied from the central nervous system.

The computational copy, he finds, is of quite poor quality. In the Fifties, computer components did indeed fail every other week, and were inordinately large for the functions they performed:

" Two reasons that put a limit on complication [of the functions of computers] in this sense have already been given. They are the large size and the limited reliability of the componentry that we must use, both of them due to the fact that we are employing materials which seem to be quite satisfactory in simpler

applications but marginal and inferior to the natural ones in this highly complex application.” (109)

The engineering challenges of building better digital computers obviously irked Von Neumann; apparently this explains why so much of his final work on automata is about the design principles for the Central Nervous System. In effect, his disenchantment with logical automata and the (ostensibly) slow pace of progress on the part of hardware engineers led him to start to envision an entirely different sort of cellular automaton. This artifact would have been constructed differently in terms of physical storage of information, the composition of the logic units, and the mathematical features of programming and control. He points out that the function of memory is indeed carried out chemically in the body and suggests doing this through biological engineering, only a few years before Watson and Crick discovered DNA (110).

The use of logics, Von Neumann supposes, will have to be replaced by some surrogate based on the nervous system’s usage of “varying chemical compositions in the blood stream “or other humoral media”. This communications and programming system he sees as being a hybrid of digital and analogue programming languages- although he does not use the latter term (111). But the means of information transmission in the CNS, Von Neumann states, is not digital but statistical. The programming and information storage hardware of a biological automaton will certainly have to involve more redundancy and tolerance of error, and indeed correction of error, than digital computers of his time could allow (112).

The study of isomorphisms between the computer and the brain, and what this could tell us about automata, was Von Neumann’s swan song. This work did not go unnoticed, instead leading toward the fields of A-Life and Cellular automata. Von Neumann’s work on these universalities is perhaps the most elegant such monograph on the topic. Von Neumann was writing the lectures that make up The Computer and the Brain, intended for the Silliman series at Yale, when he was stricken with a

dramatic case of cancer in August 1955 (113). The lectures were never delivered; he died in February 1957 (114).

3. Alan Turing's Enigmas

In the history of computing, Turing plays Dostoevsky to John Von Neumann's Tolstoy. These two towering figures could not have been more different as human beings- Von Neumann a legendary party-thrasher, a bon vivante, professionally inquisitive and networked to the hilt, Turing formal, reclusive and secretive- no surprise given that his sexuality was called a crime. It is difficult to not consider Alan Turing in a separate sphere from his fellow intellectuals, even though he was known to and worked with many of his contemporaries. His stature is pristine and sparse, rather than enmeshed with the complications that customarily take the sheen off people. His role was to pose conundrums, which he did as often as he affirmed new possibilities. Like Von Neumann, he died young (1954), unlike Von Neumann, most probably by his own hand (115). The tragedy underlines the somber tone of his proposal of what were at the time impossible feats.

Curiously, Turing did for AI what he had done for automatic computing machinery. He specified criteria for thinking machines- his term for AI- before the gauntlet had been taken up by anyone else. As we saw earlier, Turing had specified the desiderata for logical automata in 1936. His article and that of Emil Post, had helped to set off the field of study of automata, which lent a theoretical boost at least to physical automata and offered increased possibilities of abstraction in digital computing. Turing conducted significant cryptographic work and designed digital computers throughout the war, as we have seen, and after it, as well as continuing to consider logical automata. After the War, Turing persisted with computers and logic and never became interested in the emulation by computers of the central nervous system (116). His work seeded AI as he persisted with logic, and began to think about machines that thought- 'thinking machines', as the verbiage went, mostly beginning in terms of logical operations. While he was cognizant of the study of neurology

amongst other members of the Ratio Club, he never took either this approach or a real interest in cognition or psychology (117).

Turing was a logician, but his contribution most relevant to AI next to the theory of logical automata was philosophical and written in prose. "Computing Machinery and Intelligence", published in the journal Mind in 1950 and widely reprinted, poses the possibility of machines which emulate cognition, and then second-guesses and refutes various objections to AI. In the Imitation Game, which Turing proffered as a means of determining whether a machine is or is not acting intelligently, an interrogator questions a person and a machine, in physical separation from both and in the form of written notes. The machine is judged to have won if it fools the interrogator into thinking it human. The structure of the game as presented is important, and helpful to AI because it makes clear that intelligence in a digital computer appears in a mechanical input-output form.

Despite its apparent commonsense appeal, the Imitation Game (or Turing Test, TT) is as irritating to most people in AI as a traffic jam (118). The inquisitive public is unduly obsessed with it, when in fact it requests of AI commonsense reasoning, which was one of the trickier challenges to AI through the 1980s. (Understanding of specialized fields or technical procedures appeared far earlier). Despite itself, the TT tends to support a biological or overly faithful robotics rendition of AI rather than the more useful functional rendition. The Turing Test is more widely discussed (if not more widely known of) in a popular understanding of AI than in professional circles. In other words, this buoy is far more visible in the shallows than in the deeper waters.

Turing's reasoning in response to the arguments against AI is perhaps still the best brief rejoinder. He addresses, and quickly puts down, a number of anxious or simply irrelevant reactions to the digital computer's progress, most notably the idea that computational intelligence is no good if it is not like human intelligence. As one might expect from a logician, he is quite serene against worry about the physical similarity of the human

nervous systems and electrical digital computers (119). Turing's article puts the more obvious and foolish reactions to rest, but also helps point toward the vast research which had to take place before AI would really get anywhere:

“ We may now consider again the point raised at the end of part 3. It was suggested tentatively that the question ‘can machines think ?’ should be replaced by “are there imaginable digital computers which would do well in the imitation game ?” (120)

This is the sort of restatement and refinement that would have to be done many times before AI could be discussed in a more meaningful way than cocktail party babble and chatter. Turing gets to the heart of the matter. The physical instantiation of computing was not the only problem- the intricacies of symbolic representation were per se immense as well.

“ The fact that Babbage's analytical engine was to be entirely mechanical will help us to rid ourselves of a superstition. Importance is often attached to the fact that modern digital computers are electrical and that the nervous systems is also electrical. Since Babbage's machine was not electrical and since all digital computers are in a sense equivalent, we see that this use of electricity cannot be of theoretical importance.” (121)

4. McCarthy and Minsky's first Approach to Automata

Minsky and McCarthy, as we saw in the third chapter, attended college during the war years, and entered graduate studies in an intellectual environment of increasing capacity and interest in the computer's potential. Both men gravitated to the intellectual high ground of abstract automata while in grad school at Princeton.

Marvin Minsky seems to have always been the smartest kid in the class, even when the class was at Harvard College. But clearly his exceptional capacities surpassed the purely mental. His tolerance for inconsistent ideas, even those which clearly clashed, served him well at Harvard in the 1940s, and thereafter.

During his college and graduate school years, Minsky sought out the most intellectually intriguing and influential people he could find. He investigated the warrens of researchers in diverse fields deep in the obscure university hallways. He spent a great deal of time talking with B.F. Skinner in his lab during college (122). Indeed, Minsky seems to have been the only one of the AI leaders to have ever talked with Skinner (123). He met the pioneer cognitive psychologist George Miller, attended Norbert Wiener's Cambridge Cybernetics meetings, and took 'Physiological Psychology', a graduate class taught by J.C.R. Licklider, at the time an acousto-psychologist teaching at Harvard, in 1948. During his years at Princeton he worked for Claude Shannon at AT&T in New Jersey for a summer, and secured approval of his thesis from Von Neumann. The result was that when he finished graduate school he was as well connected as a concierge, and very well versed in the intellectual movements of his time.

This familiarity with the intellectual surroundings did not segue into acceptance of everything around him. As a Harvard senior, Minsky was stunned at the absence of a model of mind among psychology researchers at the time:

" There was something terrifying about this clash of two different worlds- the physiological and the behavioral. There were no psychoanalytically - oriented people around them. If there had been the situation would have been worse, not better. I could not fathom how these people could live down there arguing about personalities, with no methodology, no idea of what to do and no real theories of what is happening deep inside the mind. I could not understand how they could proceed with no model of what was going on in the mind..." (124)

Minsky and fellow grad student Dean Edmonds constructed an indigenous model, quite closely along the lines of positive reinforcement by inhibitory or excitatory impulses propagated to neurons. His naive model of information processing simulated the mind as a collection of local neurons, with binary states, which could be 'taught' to fire properly at a higher level through

repeated trials. As Minsky and his colleague Seymour Papert later described it,

“ Perhaps the first reinforcement based network learning system was a machine built by Minsky in 1951. It consisted of 40 electronic units, interconnected by a network of links each of which had an adjustable probability of receiving activation signals and then transmitting them to other units. It learned by means of a reinforcement process in which each positive or negative judgment about the machines behavior was translated into a small change (of corresponding magnitude and sign) in the probabilities associated with whichever connections had recently transmitted signal. The 1950s saw many other systems that exploited simple forms of learning and this led to a professional specialty called adaptive control.” (125)

The Snarc, as it was called for no particular reason, included several hundred ‘neurons’ meant to stand in place of undifferentiated cells with binary states, each simulated by a vacuum tube. Minsky and Edmonds built the machine in 1951 with an Office of Naval Research grant care of George Miller, and included in addition to the vacuum tubes such other parts as an autopilot from a B-24 bomber and a bicycle chain (126). Meanwhile, unbeknownst to Minsky, neurologist Donald Hebb had come up with an architecturally similar model, and published it in *The Organization of Behavior* in 1949.

Minsky graduated from Princeton in 1954, gaining the approval of John Von Neumann on a doctoral thesis entitled “Neural nets and the brain model problem”. During the mid-1950s he returned to Cambridge as a Harvard Fellow. His paper for the Automata conference organized by McCarthy and Shannon, “Some universal elements for finite automata,” bears the fingerprints of Von Neumann and McCulloch’s influence (127). Minsky proposes a cellular automaton composed from “a small number of basic elements”, that is, from something congruent to the neural system. This would be skeletally relatively simple, or so they thought. The automaton is clearly “lifelike” in that it emulates the CNS’ propagation of information in the form of Boolean pulses, rather than lifelike in its emulation of cognition (128). Minsky thus

began his career as one of the next generation of the mathematically-inclined Cyberneticists.

5. McCarthy begins to Design Languages

John McCarthy's stubbornness and independence, combined with his great intelligence, should have been good predictors of further achievement. As a senior at the California Institute of Technology, he attended the Hixon Seminar on Brain Mechanisms and Behavior. He was not interested in the nervous system, the topic of most of the papers, but went to one talk about intelligence in machines. After the seminar, he sketched out an experimental project in which rudimentary self-sustaining automata (i.e., cellular automata) would act as protagonist (learning agent) and as the environment in which the agent learns. He sent a letter about these ideas to Von Neumann, and received a reply suggesting that he "write it up." He never did this (129). But how could he have, since at the time he, and everyone else, lacked access to a digital computer ?

McCarthy's work in organizing a conference on "Automata Studies" in 1952 indicates his increasing participation in the tradition of automata. "The Inversion of Functions Defined by Turing Machines," his paper in this volume, posits the computability of a wide range of calculations, representing scientific reasoning, by Turing machines (130). The paper indicates McCarthy's longstanding (and ongoing) concern with the formalization of intellectual artifacts for computing purposes. Like Minsky, with whom McCarthy spoke often during graduate school, McCarthy was convinced that learning machines were possible, and seems to have initially approached them in the vein of physical automata. When he and Minsky worked for Claude Shannon at Bell Labs in 1953, the two searched for automata and roboticists, and found both (131). As we saw a few pages ago, Minsky had been his own robotics teacher even earlier.

McCarthy took a teaching position at Dartmouth College in 1954, and remained there through 1958. He taught during the academic year, interspersing this with summers working with Nathaniel Rochester at IBM (1955) and the long conference on AI

at Dartmouth (1956). IBM needed McCarthy's talents as a computer man to perform the programmer's beta-testing, before it distributed the 704 model computer to various regional universities. McCarthy convinced Rochester, the chief designer of the 704, that AI was work looking into. Subsequently, Rochester agreed to act as the senior member of the organizing committee for the Dartmouth Conference a year later (132). During the summer of 1955, Allen Newell visited McCarthy at IBM, and the two discussed computer languages at length (133). This interest subsequently manifested itself in McCarthy's work in computer languages, soon afterward.

As John McCarthy approached the Dartmouth conference, he appears to have been thinking of AI as a matter of automata, rather than in the psychological vein. In his 1956 paper on automata, "The Inversion of Functions Defined by Turing Machines," McCarthy declares,

" Consider the problem of designing a machine to solve well-defined intellectual problems. We call a problem well-defined if there is a test which can be applied to a proposed solution. In case the proposed solution is a solution, the test must confirm this in a finite number of steps. If the proposed solution is not correct, we may either require that the test indicate this in a finite number of steps or else allow it to go on indefinitely. Since any test may be regarded as being performed by a Turing machine this means that well-defined intellectual problems may be regarded as those of inverting functions and partial functions defined by Turing machines" (134).

McCarthy points out that a Turing machine will compute the value of a given function by putting through successive integers. But this is not good enough because if the function has no solution, the machine will not know to stop (135). In his work in the late 1950s, McCarthy would turn to the issue of figuring out the design of computer languages that could order the machine to stop.

The automata approach often devolved into physical artifacts. Yet in this work both Minsky and McCarthy were open to the

possibility of automata as something more inclusive than their original conception. In McCarthy's case, concern with the design for automata turned toward AI's defining computer language. Minsky's study of automata turned toward heuristics for the reiteration of geometric proofs in the mid-1950s, and his innumerable other contributions to AI following that.

Chapter 6. Games and the Idea of Biology as Intelligence in the Postwar Years

1. Introduction

Games of skill, such as Go or chess or bridge, with a small set of highly parsimonious rules, dictate that the winner is indeed intelligent. This is particularly the case when chance has no place. Figuring out the psychology of strategy in chess, and the ways in which a machine or a software program might play games of skill, were preeminent concerns in the early 1950s and were crucial to the formation of AI. Pure games of skill were a point of departure for researchers in both AI and automata because they offered a small, closed world in which no other rules or facts needed to be referred to. Problem-solving could be approached very simply and schematically, leading to early modeling of problem-solving and search itself.

2. Von Neumann and Morgenstern at Rand: Games No One Wins

Widespread fascination with chess, checkers, and other games of skill in the mid-Twentieth century was perhaps the mark of a society which was less mechanized, and poorer in expensive leisure activities, if richer in leisure time. The Depression, for instance, was the Golden Era of board games, card games, miscellaneous contests, and of course illegal gambling. But as everyone who has watched a game played for money knows, games can be more than games. The stakes may exceed even money: strategies for playing war games involve nothing less than life or death.

Following the war, there was less leisure and more at stake: games became an academic topic rather than simply a leisure pursuit. The spirit of skilled game-playing became deadly and

desperate rather than one of idle entertainment. John Von Neumann and Oskar Morgenstern published The theory of games and economic behavior in 1944 (136). This monograph instigated the study of games, re-interpreted as simple models of economic behavior, in mathematics, economics, war strategy, and Cybernetics. In this new genre of study, the world is reduced to a schematic forum in which nations or firms contend. For instance, given that the goal state is maximizing income, and that the moves of the opposing firm in a two-player market are hidden, the protagonist firm may pre-emptively enter a new market, hoping to capture consumers' loyalty. Game theory became a well-traversed micro-economic sub-field.

Another intellectual sport was the application of games to nuclear war. Unlike traditional games or using games to represent the competition of firms, this usage of the concept of games was preposterously unrealistic and far from innocuous. This was played at the Rand Corporation throughout the Cold War decades. The sheer oddity of some of the work done at Rand was not lost on its denizens:

“ The RAND Corporation's the boon of the world
They think all day long for a fee
They sit and play games about going up in flames
For counters they use you and me.’
“ The RAND Hymn,” by Malvina Reynolds. (137).

The reduction of the spectre of death and destruction to the arcane abstractions of a game is a good example of finding a euphemistic way to thinking about things one can't really bear to think about. The terms 'nuclear' and 'war' do not dovetail nicely, or rather, honestly. War as the terms is historically used is not even the proper word for describing any sort of contention involving nuclear weaponry. Traditionally, a war continues for a protracted period marked by intermittent altercations. In contrast, nuclear war is 'won' or 'lost' in an instant. But through such willful mis-statements about nuclear war as a game (138), the nuclear 'game' or pseudogame continued for decades (139).

3. Chess and Checkers as a Sequeto Problem-Solving

Whatever the putative topics of game theory, this and increased highbrow attention to actual board games had salutary effects on the self-esteem, so to speak, of Cybernetics. Investigating games implies such active, rather than passive, things as figuring out strategies, imagining offense and defense, considering long-term versus short-term moves, and the like. Physiological psychology is rather passive by nature regarding the contents of human thought, but the investigation of game-playing is by nature the opposite. Because of their popularity in the English-speaking world, checkers and Chess were the most notable of games which became popular in early work on computer game-playing in the 1950s. The field of game-playing itself epitomized the transition between computer programming which saw human intelligence as a form of conditioned reflex productions, versus one which emphasized cogitation. In the early 1950s, chess-playing and checkers-playing's adherents took up information-processing views of intelligence- ersatz in the case of Shannon and Samuel's, and explicit and evangelical in the case of Newell, Shaw, and Simon.

Claude Shannon sets the Bar for Chess Analysis

It is said in AI that chess was the drosophila of the field. This was the case even before AI existed. Indeed, chess was the drosophila both as the effect of the formation of AI, but also its cause. The intense demands of the field helped to give rise to AI, since examination of the game repeatedly showed that chess posed too great a challenge to blind search, certainly, but also to any form of search which did not try to embody problem-solving strategies. Inevitably, asking such questions evoked the issue of human problem-solving. The game appeared and reappeared, and continues to do so. The definitive presentation of the problems of chess, if not of its solutions, was and perhaps remains, that of Claude Shannon, in "A Chess-Playing Machine" in Scientific American (1950). Another equally weighty article appeared in Philosophical Magazine in 1949; both of these are referred to today for their clarity as well as their historical value.

Shannon points out that electronic digital computers are powerful enough to carry out symbolic computing, rather than only process numbers, then indicates how the squares of a chessboard may be expressed in machine code (pseudocode, actually). He indicates immediately that a blind search is not a possibility: chess just involves too much data. A machine engaged in chess using blind search, in fact, would have to take on 10 to the 120 possible moves. Chess almost invariably forces the thinker toward strategy, and Shannon suggests a numerical evaluation score to indicate how the player(s) are doing, with respect to certain obvious features of chess-playing. Then he turns to Dutch psychologist de Groot's studies of the strategies of expert players, and suggests subprograms to test out various plays to several ply. In the Philosophical Magazine paper, Shannon proposed in a more technical manner the form he expected problem-solving to take. The finite and non-deterministic nature of chess dictates that the game may be construed as a branching tree, albeit always portrayed upside down, from which springs nodes (the positions) and branches (the moves).

The strategic portrayal of the sequences of a game may be pursued through such a graph, with the maximum benefit being chosen by the computer's side being played, and the minimum benefit being chosen, naturally, by the opposing side. The alternating plys are referred to as mini-max because of their respectively differing objectives (140). Shannon's paper unequivocally showed up the inadequacy of the existing neurological or conditioned reflex model of information processing, and began to suggest alternatives. Widely read and highly inspirational, his paper helped to instigate the introduction of cognitive modeling (141).

Shannon's papers were a catalyst for various chess-playing programs in the mid-1950s. The work was done by a number of groups in different research centers, a grassroots movement of haut scientific intellect. Independently of each other but not of Shannon, these elaborated further concepts in strategies of

game-playing. All of which research underlined the feasibility of formal methods of game-playing. Throughout the 1950s and even into the 1960s, practically every AI practitioner was an aficionado of computer chess-playing as well. Alan Turing's proposal for a chess-playing program to run on the MADAM computer, worked out during 1947 and 1948, implemented the concept of a "dead" position, which is relatively stable for several moves further and can thus be evaluated (142). A group of physicists at Los Alamos National Laboratory programmed the MANIAC I computer to play chess on a 6x6 board, with some pieces removed. Their program did not keep records, which is obviously essential for learning, but used a polynomial linear equation to evaluate moves, as well as using a minimax procedure. This program could beat beginners (143).

Another contemporary program that merits attention is that of engineer Alex Bernstein, who began computer chess playing independently while at the Bureau of Standards. He developed his program amidst the publicity of IBM's too-famous New York Service Bureau. This program, which NSS considered very important, ran on the IBM 704 and used an 8 x 8 board. It employed subroutines, which generated prospective moves, each substantively related to the stuff of winning a chess game. For instance, there were generators for defending the king, attacking the others' pieces, etc. While using the other features such as the dead position and static evaluation, it also introduced an efficient dose of pruning prior to search (144).

Computer scientists were not the only people interested in computer chess-playing: the general public was too. Bernstein's program was featured in Scientific American, the New York Times and Life magazine (145). While today computers and their most exotic applications are considered to be of great fascination, such was not the case sixty years ago. It is certainly indicative of the public image that IBM wished to portray that IBM, like Queen Victoria, was not amused:

"So here was Bernstein getting what for Watson was unpleasant notoriety about his chess-playing machine and Arthur

Samuel, reaping a harvest of publicity for his checkers playing program. Faced with these two very successful examples of homegrown AI- soon to be joined by a third, Herbert Gelernter's geometry theorem proving program- sales executives at IBM began to grow nervous lest the very machines they were trying to sell prove so psychologically threatening that customers would refuse to buy them. Thus they made a deliberate decision to defuse the potency of such programs by conducting a hard sell campaign picturing the computer as nothing more than a quick moron" (146).

The Bernstein program was clearly more elaborate than the other programs on the topic, and each subsequent program was more extensive than the previous ones. But, to be fair to them, we must note that their work literally predated software- the landmark pioneering text had only appeared in 1951 (147). Moreover, the Los Alamos program was simply offered more scanty resources. It is not surprising that these programs employed strategy without trying to figure out how a human being would or did play chess. But it is equally unsurprising that when Newell et al. took up this topic they looked at De Groot as well.

4. Arthur Samuel's Checkmate in Checkers

One program which contributed to AI was a much humbler game than chess, but nevertheless a fine test-bed. The odd field of Checkers is epitomized by the mixed experience of Arthur Samuel, who became known as the world's greatest expert in checkers-playing software. This title, as we shall see, is what right-handed people call a left-handed compliment. Samuel, an electrical engineer with expertise in vacuum tube design and a taste for digital computers, moved from Bell Telephone Labs to the management of an electrical engineering laboratory at the University of Illinois in 1946. He began lobbying on his own behalf- namely, for funds which could help him to build a digital computer. But these could not be bought for any price at that time. Samuel asked both the Eckert and Mauchley team and the IAS team, but neither was ready to offer one (148). If you wanted one done, you had to do it yourself.

The new engineering school Dean, the business-minded Louis Ridenour, procured \$110,000, even more than they had asked for (149). It still was not enough money, and they still could not finish on time the small computer which they could build. But here Samuel found his claim to fame, or rather it found him. While trying to use the small computer the team was building as leverage for funding for a larger computer, Samuels decided to show off the computer's checkers-playing ability, as checkers appeared at first glance to be a trivial game (150). But appearances deceive: Samuel thought checkers was humble, but it turned out to be uppity. Checkers appears very easy, and so is accorded low prestige. But in truth, according to Samuel, it is a difficult field of study.

Samuel soon became more interested in actually designing a computer at an institution which could afford to build one properly, namely IBM Poughkeepsie. IBM, knowing of his expertise in vacuum tubes, hired him early in the 1950s because he knew about the technology. Instead of endorsing the tubes, however, he proved instrumental in persuading them to adopt transistorized logic units sooner rather than later (151). Consequently, the later commercial versions of the Defense Calculator (704) used transistors, while only the early 701 and 702 used vacuum tubes. While designing the 701, he used the checkers program as the testbed, and thus some of the earliest software, progressively turned from pseudocode to assembly and machine code (152). Samuel also took a clue from Christopher Strachey of Manchester University, who had programmed a Ferranti Mark I computer to play checkers (or draughts, as the English insist on calling it). Strachey's presentation of his program at a North American conference prompted Samuel to use checkers (153).

Samuel proceeded with the program on the 701 in 1952, then developed a program with learning, which appeared on television early in 1956 (154). His programs, developed in successive formats between 1952 and 1960, were not hampered by the fantastic combinatorial explosion which characterized chess. Therefore looking ahead could take place without pruning, the machine could handle the volume of data that the prospective game plys offered. The programs, interestingly enough, were

quite similar to the chess playing programs. Samuel offered a linear polynomial which could evaluate the current position in light of various desiderata. As Samuel points out, one cannot simply assume the best strategy from the side one is playing, “since to reach this position would require the cooperation of the opponent.” Instead the strategy must be to carry the minimax procedure backwards to select a move.

The Checkers program, which eventually beat top checkers players, was exactly what the mass media wanted to see. It was so photogenic that it embarrassed IBM’s cultivated image of square, stodgy dependability. IBM moved Samuel to Europe to do research on the computer industry there, to reduce undue publicity. IBM regularly did this to preserve its subdued public image (155). This backfired because it garnered even more attention for the program. Meanwhile, Samuel never learned to like checkers and “could not really learn anything from the literature”. His work on the program, and the work of others, proceeded, though. The Checkers program reached the masters level by 1961- although Samuel himself never reached this level (156). After retiring, he visited the Stanford Artificial Intelligence Lab as a domain expert- and could not develop an expert system to play checkers.

Despite the toy problem of checkers, Samuel proved influential in clarifying game strategy. During the early 1950s in particular, his prominence in computer models of games made him influential to the founders of AI. He attended the Dartmouth Conference, and contributed to heuristic problem solving and game playing throughout the early years of AI.

5. Allen Newell Dabbles in Chess

One founder, Allen Newell, appears to have made his first embryonic statement of AI in a paper considering the chess problem. In his first year at Rand, Newell became interested in the problem of chess during the year after he had decided that he would pursue problem-solving through cognitive emulation, but before he and Simon had taken up the strategy of the Logic Theorist. Shaw indicates that he and Newell began discussing this topic at length, which would eventually lead to Newell’s paper and to their work with Simon:

“ In 1954 I heard Al Newell walking down the hall talking about chess. And that was the beginning of my interest in getting together with him to devise a chess-playing program.” (157).

Newell and Shaw discussed chess, and this resulted in a chess-playing program and in Newell’s paper on the topic. The chess paper is an intellectual mishmash- albeit one obviously written by an exceptionally smart person. Newell muses his way through the sequence of prerequisites already traversed by Shannon- modern computers are highly flexible, blind or exhaustive search in chess is not feasible, static evaluation of the player’s position must be carried out but is difficult without taking on the substantive matter of particular strategies. This is all true but simply a thoughtful restatement of things said by others. But Newell then turns from this essay away from some of the obvious and expected suggestions. Instead of working out a minimax truncated version of chess, he turns toward heuristics. Certainly this is meaningful and is Newell’s first published indication of this orientation.

Given his intellectual pedigree and his work at the time for Oscar Morgenstern at Rand, a game-playing approach would indicate that Newell was falling into place in his environment. But Newell did not fall into place by suggesting chess as a form of game-playing as sanctified by game theory (158). Instead he offers the option of using rules (the term heuristics does not yet appear). He suggests the usage of subgoals, or rules of thumb, which may be manipulated in the way that logical expressions are, and which ones will be rewarded (“reinforced”) when they succeed (159). This is thus a rudimentary suggestion of a learning mechanism for chess, couched in the language of the contemporary intellectual milieu. Newell started thinking about this in 1954, before his perspective may be strictly referred to as an AI one. Newell, Shaw and Simon began computer chess-playing research in 1956, and published their own program, which we shall examine in Chapter 8, in 1958.

As we will see in the next chapter, the work on games and other intelligent tasks began drawing numerous people toward an explicit statement of AI itself as a goal. During the early 1950s, this had not yet been clarified. Although today it appears like one

obvious task for computing research, at the time it took years to be made clear.

Chapter 7. The Impasse at the End of the Cybernetic Era

Introduction

As the Fifties commenced, there were numerous different means of discussing and trying to construct intelligent artifacts. Metaphors of intelligence still focused on the transmission of information as a sort of railway station interchange rather than its qualitative processing. Cognitive psychology had been initiated by the early 1950s, yet it was in its infancy. It could not inform any interesting proto-computer researcher of the nature of problem-solving, memory, or the structure of creativity.

Stirrings of ideas about the outer limits of computing- in the form of 'thinking machines' as Turing put it, or complex information processing or AI- could not yet be aided by psychology. The idea of a program that presented and processed a declarative computer language for problem-solving was not yet consolidated. This sounds circular, but theoretical concepts and tools to carry out various functions inspire each other.

Psychologists did not yet undertake the direct route of understanding human intelligence as a matter of problem-solving. Instead, they used Behaviorist models of conditioning or the Cybernetic component of feedback in order to understand learning. These were intrinsically indirect approaches: human problem solving would have to be understood using clinical psychological testing and using information processing as a model (160).

Thus the intellectual world in the realm of AI and its fraternal disciplines of computer science and cognitive psychology were between paradigms as well. By its very nature, Cybernetics, the idea of a meta-science devoted to terminology used to describe intelligent organic and inorganic artifacts, would not provide a research agenda for basic work in research disciplines. This had never been its stated goal to begin with, as it had been intended to provide a high-level terminology for discussing traits of intelligence in both machinery and humans. In addition it suffered

from increasing internal schisms (161). Cybernetics slowly waned in spite of its surfeit of intellectual horsepower, for lack of an overall direction. This was a most fruitful decline, however, as the numerous modern sciences of the mind were nurtured by its attention to the abstract concept of intelligence.

Cybernetics, automata, discussions of Universal Turing machines, all began to establish a way to make sense of intelligent machinery, even as metaphors for its conception were unclear, and its applications still were still in their infancy. Even as it was difficult to know what to make of something so obviously protean, the computer was establishing a permanent but not well-defined niche in American cultural life.

2. The Popular Understanding of the Computer as a Giant Brain

The announcement of AI as a field did not take place until 1956. However, as we saw there were numerous other ways of approximating intelligence in machines. In the previous chapter, we saw that both Von Neumann and Turing approached the concept of automata, and that both men were concerned with automata conceived in a highly abstract form, as basically forms of software. However, the discussion of automata did not use the terminology of intelligence or “thinking”.

No matter- once digital computers had been created, the idea of computers ‘thinking’- whatever that might mean- caught the increasingly broad public imagination. The idea of ‘giant brains’ is on its face frightening as well as patently technically muddled. The very title of Edmund Berkeley’s 1949 book Giant brains; or, Machines that Think, one of the most popular scientific visions of the time, speaks for itself. Suggestions as to the uses of machines put together many applications, which at this point are clearly ramifications of entirely different technologies. Berkeley suggests that we will have ‘automatic’ address books, cooking machines (operated by program tapes), speech recognition and handwriting recognition capacities (162). These are all indeed possible now, but are the products of different lines of research, which took place at different speeds. Voice recognition and synthetic speech appeared in commercial and feasible form considerably later than word-processed files.

Berkeley and other proponents of computing discussed how giant brains would enrich our lives. But this was often more well-intentioned than well-defined: Berkeley, for instance, is not too far from a popular and psychology textbook misconception when he suggests that the reader think of a mechanical brain as a sort of railroad station, with input, storage, a computer for information processing (with no detail on the computer) and output, and a control function and control line (163).

As we saw in Chapter Two, this pop psychology metaphor persisted from the Twenties through the end of the 1940s.

In his defense, Edmund Berkeley was optimistic. Cultural premonitions significantly exceeded the field's grasp at the moment- and did specifically address AI- or thinking machines, since the term did not exist yet. The anticipations, not all of them benign, were dead-on in intimating the importance of the field.

3. Early Attacks on AI, and Alan Turing's Response

Anti-machine sentiments are as constant an historical theme as is the desire to build machinery that embodies intelligence. The ongoing animosity of efforts to embody human abilities in machines and resistance to those efforts is part of modern culture. It is seen in Mary Shelley's Frankenstein (1818) and in the machine-breakers of utopian novelist Samuel Butler's Erewhon (1872). As they took on greater aspirations as to what they could do, they would provoke more controversy. In addition to free-floating speculation, talk about computers appeared, much of it mealy-mouthed. Even at this early date, humanists were railing against the anthropomorphism of the machine, while simultaneously staking out the ground in advance by asserting that digital computing machines could do nothing that was human.

Because of the general dearth of understanding of problem-solving, memory, speech utterances or other cognitive acts, several pernicious fallacies about thinking arose during the late 1940s and early 1950s.

The British Museum algorithm, coined in the few years after the war by a prominent but skeptical scientist, Sir Arthur Eddington, is a justification of the falsehood that intelligent acts arise out of blind search. It refers to the unlikely instance that a

troop of monkeys given typewriters would, through sheer chance, eventually type out the book inventory of the British Museum. Heaping impossibility upon impossibility, this scenario is also known as the “fifty million monkeys technique” (164). In current AI parlance, the phrase blind search indicates a traversal of a search space- that is the world of all the possibilities to a goal- without any knowledge or strategy.

In Faster than Thought (1953), B.V. Bowden inventoried tasks which computers could and could not do. A mechanical brain could, he said, learn what it is told, remember numbers and perform arithmetic, find numbers in tables, and write out and verify numbers, among other things. But machines could not, said Bowden, think intuitively, leap to conclusions, or interpret ‘complex situations outside itself’. In the field of learning, Bowden says, electronic digital computing machines could not reach beyond conditioned reflex (165).

This was not helpful, either. Both the statements of limitations and those of capabilities ignore input and output, computer languages, applications programs, problem-solving in computers, memory and the restriction on applications imposed by its terribly small size, natural and formal languages, etc. The only solid thing indicated in fact is the well-established ability of computers to perform arithmetic operations. The vagueness of this pro-computer statement let open the way for coffee-shop fly-swatting sorts of attacks on AI, such as the following one cited by Bowden:

“ In his Lister Oration for 1949, Professor Jefferson said: “Not until a machine can write a sonnet or compose a concerto because of thoughts and emotions felt and not by the chance fall of symbols, could we agree that machine equals brain- that is, not only write it but know that it had written it. No mechanism could feel (and not merely artificially signal, an easy contrivance) pleasure at its successes, grief when its valves fuse, be warmed by flattery, be made miserable by its mistakes, be charmed by sex, be angry or depressed when it cannot get what it wants.” (166)

As it would do later, AI acted as a lint brush for the ambient anxiety of the age. This is quite ironic, given that prior to 1956, the name and the practice simply did not exist. The general sense of unease was certainly justified, even if the target was the wrong

one. As the Cold War's outlines sharpened, it became clear that no one intended to lay down their arms and beat their swords into ploughshares. For anyone already anxious about the breakneck development of nuclear weaponry, the idea of thinking machines, prior to the existence of the term "AI", must have seemed a likewise appropriate target of nervousness. The transmogrification of existing nuclear weapons and computers into a terrible beast that embodied both must have seemed possible.

The legitimacy of the fear does not mean that AI was indeed the lackey of nuclear warriors. It was instead one of the earliest pacifistic and scientific purposes to which computers were turned. As we will see in the next chapter, Newell, Shaw, and Simon met and began working together as a result of the SAGE anti-missile defense project. Neither any one of these scientists nor the other AI founders were either politically active or deeply engaged with the Cold War. However, the Cold War ruthlessly drew in every available technology, including rapidly developing computer power, and AI's founders had to be involved in some way in order to pursue their research.

Objections to Artificial Intelligence thus preceded the field's announcement. In 1950, Alan Turing's "Computing Machinery and Intelligence" indicated that the general idea that machines would be able to 'think' someday was often discussed, and that it made humanists anxious (167). The digital computer had just been invented, was scarcely functional, and so criticisms would appear to have been yet unwarranted. Turing's strike was preemptive, yet clearly aimed at tangible objections that he had heard.

Turing's paper "Computing Machinery and Intelligence" dissects various arguments objecting to computer simulations of intelligence. While published well before widespread usage of commercial hardware or software, its distilled arguments would persist for decades (168).

Some of these contentions are unilateral statements with little meat on the bone. Turing points out arguments by stipulation, that is ones which simply assert, "Let's settle this once and for all, machines cannot think!" or "A computer is not a giant brain... It is a remarkably fast and phenomenally accurate moron". Another ephemeral argument is the theological one: 'Thinking is a function

of man's immortal soul.'(169) The 'heads in the sand argument'- 'The consequences of machines thinking would be too dreadful. Let us hope and believe that they cannot do so'- is futile as well as lacking in any attempt to refute assertions against it.

Another attempt to simply define AI out of existence by asserted that intelligence per se requires biological life as humans know it (170). This lacks bite because it does not even attempt to address head-on the tasks that computers may actually undertake.

Some of the arguments against AI are really arguments against computers and computing, and are evidence of the general level of ignorance found even in the educated publications of the day. Usually the arguments assert that computers cannot remember enough material to be useful. This Turing lets the proponents of the early anti-AI arguments off the hook, so to speak, by pointing out that these ways of thinking are:

"...mostly founded on the principle of scientific induction. A man has seen thousands of machines in his lifetime. From what he sees of them he draws a number of general conclusions. They are ugly, each is designed for a very limited purpose..." (171)

'Lady Lovelace's objection' (1842), found in her notes to the de Menabrea lecture transcription, asserts that, "The Analytical Engine has no pretensions to originate anything. It can do whatever we know how to order it to perform (her italics)." (172) Turing points out that this is practically not really true for a rather mundane reason:

"Machines take me by surprise with great frequency. This is largely because I do not do sufficient calculation to decide what to expect them to do..." (173)

More sophisticated arguments would appear later, featuring a more precise idea of what is meant by thinking itself (174). None of these is particularly compelling as given in Turing's early work. However, as we will see in Chapter 12, none of this stopped AI's detractors.

4. The Pregnant Pause in the Interregnum

We have defined the Cybernetic interlude as a distinct period, nearly coterminous with the Postwar years and the Fifties. It was also the immediate predecessor to AI. Characterized by

fascination with gadgetry, uncertainty as to the nature of the relation between machines and people, vague characterizations of computers as 'giant brains', and thoughts of intelligence as a form of telephone switchboard, this period could not have spawned AI, and did not. In retrospect, the digital computer's growth took place rapidly. But in actual life it happened fairly slowly, despite the growing ranks of brilliant minds that were devoted to it and the staggering innovations that took place.

In 1953, Bowden refers to "the remarkable device often called electronic brains", and opts to call them 'electronic digital computers', rather than assuming that the term speaks for itself (175). But, then, the term does not speak for itself. There were many versions of computers in use, with a handful of different hardware and memory forms. In Berkeley's book, many different types of machines are presented as viable alternatives to distinct sorts of problems- calculating punches (which were still being invented in 1946), analog differential analyzers, the Bell relay calculators, and even the Kalin-Burkhardt logical truth calculator and harmonic analyzers, are presented as roughly equal contenders for different roles. The heterogeneous presentation leads to an impression of more enthusiasm than clarity. Berkeley's book sees many types of computers as equal, when in fact history shows us that the electronic digital computer was the homo sapiens among the Cro-Magnons and Neanderthals.

The emphasis upon the distinct purposes of different forms of hardware, rather than a uniform computer and the distinct software that could be designed to run upon it, makes the taste of the literature of the time quite foreign. The emphasis upon distinct dedicated machines rather than software certainly fed the Cybernetic train of thought, which emphasized the machines' emulation of biology rather than the later AI approach to machines as 'thinking' entities. The biological approach persisted for the duration of the period through 1960, at which point AI and cognitive science were much more firmly underway.

During the 1940s, efforts at trying to envision the mind as an information processing device failed in part because of a dearth of terminology. Pamela McCorduck indicates that improvements to the field psychology took place by stealth even during the Forties:

“ Herbert A. Simon has brought to my attention the presidential address before the Eastern Psychological Association by Edwin G. Boring in 1946. While Boring has no notion of the information processing theory of modeling the mechanisms of mind, he reports a discussion with Wiener who “defied me to describe a capacity of the human brain which he could not duplicate with electronic devices. I could not at once name him any, and I confess that I myself thought it would be salutary to show that all human mental functions can have their electronic analogues. I lacked, however, an inventory of the functions, and thus could not be sure that there was not some psychological function left over, one which a nervous system could perform and an electronic system could not.” (176)

Here Boring refers to the notion of search as an intelligent procedure, a vague notion of what Simon now calls the chunking process (the ways in which memory associations are made), a sense of an information processing view of the nature of symbols, and a proposal for a Turing-type Test four years before Turing is to suggest it. A study of robotics, he says, will force psychologists to stop using vague terminology now incapable of rigorous definition. Boring and one or two other psychologists were in the minority; the dominant paradigm was stimulus response.

As the Fifties began and progressed, a surfeit of study of different sorts of theoretical and real intelligent machines approached AI without actually stating it outright. This may be in part simply because of the novelty of the computer.

However, the digital computer was not yet the singular computing genre for information processing. Cybernetics abounded and the emulation of other sorts of information or intelligence bloomed. The biological metaphors of Cybernetics also certainly drew intellectual vigor away from the potential for thinking of software as a metaphor for thought. The period earlier had encouraged looking at all sorts of intelligence because of cognition, almost naturally, was thought of as a neurological, almost telephone-like process of getting the calls switched properly. The multiple flora and fauna of this period would remain preeminent until the digital computer worked much better. Instead of understanding thinking as a protocol or other means by which to approach thinking. Because there were no computers by

means of which to run computer software programs, there was no way of practically demonstrating theories of intelligence using computers. Such theories could not be based on cognitive psychology experiments because the academic discipline did not exist. Once computer input could be turned into output in a relatively fast manner, and once computer languages and systems removed at least some of the tedium from writing programs, computational emulation would become more attractive than direct robotic-type emulation of various human I/O. Then there would be emphasis on more parsimonious methods (that is, the more restricted I/O of computer software) and naturally more interest in thinking. The work done under the intellectual banner of Cybernetics would continue as robotics and the like. The pecking order of pre-eminence would shift somewhat as the bright young man (or, occasionally, woman) heading into computer science fields became much more likely to head into AI than robotics. Eventually, however, these roads converged- but that is a story this author will tell later.

How did the well-articulated and indeed global gestalten of Cybernetics give way to the very different approach of cognitive science and its computational counterpart, AI ? Some of the most brilliant and open minds of the time broached on the very turf of AI without ever seeing it. It is not in doubt that many of those who later entered AI started out in the Cybernetics realm- Minsky indicates that McCulloch and Pitts were quite influential indeed (177). As we saw earlier, Cybernetics was indeed entering a period of slow explosion and dissipation, but this fate was by no means clear in 1955. Why AI showed up as it did, in its 'cognitive' or 'Classical' form, so early, was a result of a number of historical wildcards. The unexpected feature in the appearance of AI was certainly the unique meeting of Newell and Simon, and various conversations between them and Cliff Shaw and Oliver Selfridge in 1952-1953.

But there were other events, generally independent, which created an historical 'open point' in which AI could start, as Allen Newell so felicitously put it. These include the dramatic improvement of digital computer memory, logic units, and I/O circa the middle to late 1950s; the appearance of rebellious young scholars in cognitive psychology and linguistics, and the

increased prestige of sophisticated emigre Gestalt and Freudian psychologists from Europe. In the broadest analysis, it is not clear that the two were really entirely independent. Instead, highly diverse events emerged, ultimately, from a single cause- so large in scope as to be useful as historical narrative and useless as prescriptive policy. The fine hand of the military industrial complex could be found working here. Cognitive science was doubtless able to get started because there was more research funding for both the establishment and the rebels in psychology. All of these things certainly hastened the emergence of AI out of the Cybernetic period.

By the early 1950s anyone questioning what intelligent processes actually were was beginning to run out of ways to answer these questions. They might indeed find themselves all dressed up with no place to go, so to speak. An impasse was reached.

The level of intellectual capacity and energy that the new generation collectively exuded was impressive. Notwithstanding this, many of their enterprises consisted of independent hobbyist robot-building, speculation as to why thinking machines were in principle possible, and high-level epistemological discussions of intelligence without much experimentation in the substrate. As we will see in the next two chapters, the best way to breach such an impasse is to cut through the Gordian knot, with little fanfare. AI would cut through this impasse, first for the work of Newell Shaw, and Simon equipped with an early digital computer, and for an entire intellectual generation with the declaration of the field itself in 1956.

Chapter 8. Newell, Shaw, and Simon at the Rand Corporation

1. Introduction: Rand as Postwar Behemoth

“ The RAND Corporation- a nonprofit corporation formed To further and promote scientific educational and charitable purposes, all for the public welfare and security of the United States of America”. Rand Articles of Incorporation. (178)

It is not only politics that make strange bedfellows; research science engenders odd meetings as well. AI, a discipline carried

out by dedicated intellectuals who rarely involved themselves in politics as such, began with a close association with both uber-Cold Warrior John Von Neumann and the Rand Corporation, an institution central to the United States-Soviet conflict. Yet perhaps there was design rather than paradox here: the weapons designs systems that the Cold War required drew in numerous academics for intellectual projects relating to the improvement of the computer. This broad agenda included ones that led to AI.

Project Rand, originally meaning Research AND Development, was formed in 1945 as a special division of the Douglas Aircraft Company, for work as an intellectual auxiliary to the Air Force. More precisely,

“ In 1946 the Project RAND objective was stated as a program of study and research on the broad subject of intercontinental warfare other than surface” to include recommendations of preferred techniques and instrumentalities” to the Army Air Forces.” (179)

Three years later, Project Rand was renamed the Rand Corporation, headquartered in Santa Monica near its parent institution. It became the gathering place for Department of Defense “power intellectuals” through at least the first two decades of its existence (180). During the 1950s, hundreds of government and university researchers converged at Rand each summer. Rand scientists speculated as to the feasibility of communications and weather observation satellites, air traffic control, the use of titanium, urban mass transportation (181), and how mankind might survive nuclear warfare (182). As Alex Abella tells us in his biography of the institution,

“ At the close of the decade... RAND had amassed the grandest collection of brainpower in one institution since the Manhattan Project” (183).

With far less press coverage, Rand also was the forcing ground for the earliest AI, done by Allen Newell, Clifford Shaw, and Herbert Simon, and for the decades-long collaboration of Newell and Simon in AI and cognitive science.

Rand was astonishingly favored not only by the Defense Department but also by various agencies and philanthropic foundations:

“the Office of the Secretary of Defense, the Atomic Energy Commission, the National Aeronautics and Space Administration, the National Institutes of

Health ... The Ford Foundation, The Rockefeller Foundation, and the Carnegie Foundation" (184).

The work that Rand undertook was unequivocally lucrative. It included the design of computers, some of it paid for by federal largesse; the subject of game theory, much beloved by the nuclear strategists; and parts of the Semi-Automatic Ground Environment (SAGE) Soviet air strike warning system project. John Von Neumann, one of our book's heroes and an unequivocal contributor to the intensity of the Cold War, summered at Rand. Oskar Morgenstern, another founder of game theory as discussed in Chapter Six, conducted research there as well (185).

In both the wealth of the institution and its basic structure, Rand differed from colleges and universities. Universities are still, and certainly were at that time, designed to instruct in basic disciplines, both broad and narrow, but not to serve expressly specific missions, or technical challenges (186).

But Rand did not do the tedious work of teaching histology, the Congress of Vienna, or differential equations. It brought those already knowledgeable in academic disciplines to look at specific challenges, and paid them better than universities. The results were often superlative. Yet the open-endedness of the mandate itself inevitably led to derisive quips:

Q: Specifically, what are some of the examples of the Center's work ?

A: Well, the Center staff members have resolved the conflict between teaching and researching.

Q: How ?

A: By doing neither". (187)

Given these lucrative grants, Rand had an exceptionally wide berth in its projects. It could keep its building open to cater to those who worked odd hours- at the time, any hours outside of nine to five, weekdays (188). It could design its facilities to encourage people to interact (189). Rand could host dozens of the nation's best mathematicians for the summer to solve defense problems. Its managers could offer researchers designated time to work on their own research. At this time, it may have seemed that Rand and its ilk, rather than universities, might be the academic wave of the future- and indeed many of these innovations have been adopted at other research institutions.

But the idea that universities would wither in favor of think tanks turned out to be incorrect. Universities endure: the think tanks were a new variety of intellectual species, rather than a new successor displacing the incumbent.

2. The System Research Laboratory Project

In addition to being a novel academic institution, or a research-academic hybrid, Rand carried out immensely novel and challenging research. One of the most notable of these programs was the Sage project, and its subsidiary System Research Laboratory or SRL.

Finding out about AI's integral relation to the SAGE project ('Sage') is similar to an adoptee's opening the books on his or her birth mother. One is not guaranteed the mother of one's dreams. But if bigger is better, Sage was a pretty good Mom to many technologies. Sage was an effort to construct a comprehensive, rapid, and accurate early warning system for the anticipated Soviet air strike against the continental United States. This mission provided ample opportunity to implement control systems that were amenable to human use. Sage had to have in place the hardware to sense an attack and coordinate a response in kind. The hardware research, as we have seen, produced such advances as timesharing, electronic mail, improved memory for storage and for programs for computing, and computer languages which could process such data (190).

This work was done at Rand and elsewhere. The man-machine interface, as it is called, was Rand's specialty, as established in its foundation as an Air Force proxy. As an experimental laboratory for the Air Force, Rand was called upon to test control systems, that is, machines that control other machines. To this end, Rand undertook a grand and expensive project known as the System Research Laboratory, at March Air Force Base, in Riverside County, seventy miles east of Santa Monica. The project aimed to gauge and improve the human operators' responses to the prospective bad news. The SRL simulated an anticipated airborne missile attack, and then assessed the military personnel's response (191). From today's safe vantage point, it appears like quite an elaborate mock-up. At the time, it was a very serious dress rehearsal:

“ In this model, the machines were radar sets and fighter planes, and the men were plotters who had to trace on the surface of a large Lucite screen the location and direction of aircraft spotted by the radar, if the craft were unknown, a decision would have to be made as to whether a fighter should be scrambled to go out and look at it. To scramble a fighter for every unidentified object was costly; to miss an enemy plane would be costlier still.” (192)

The pseudo-attacks appeared realistically on the radar team displays. At its peak, the SRL was highly authentic, with “deep simulations“. The SRL staff at first used college students to react to the simulated air attack. The students had to be replaced, as they took naps except when a crisis was unfolding (193). Later, they brought in an actual military crew. It was discovered that the operators began to filter information above a certain threshold, a result that was meaningful for both empirical cognitive psychology and for the display of information in technical environments. The scale and aplomb of the Systems Research Laboratory was itself indicative of the expansive mood of military science at the time. In his autobiography, Herbert Simon calls it:

“ a grandiose project if there ever was one outside physics and space science...Al and his three colleagues simply took it for granted that it was reasonable for the Air Force to build an entire simulated air defense station and to staff it for years with an Air Force unit, enlisted men, and officers.” (194)

Both Newell and Simon were involved in the SRL, as we shall see in a moment. Simon thought the System Research Laboratory was far-fetched, but Newell was enthusiastic. Regardless, its ongoing practice garnered the repeated presence of both Newell and Simon at Rand through the entirety of the Fifties. They would conduct other projects there early in the 1960s. Moreover, the SRL, with both Newell and Simon as active participants, contributed mightily to the formation of the idea of complex information processing- their term for AI (195).

3. Allen Newell Finds his Metier

Simon was already finding that human decision-making was at least as interesting as theoretical economics. The SRL helped the much younger and less well-defined Allen Newell find his way. At

the start of the 1950s, Allen Newell was younger and less well-articulated as to his planned prospects than were the other leaders of AI. He had not earned a graduate degree, the universal prerequisite for teaching. Yet the lack of dedication to one topic or any career path at the turn of the 1950s had turned into serious dedication to AI by the middle of the 1950s. Indeed, this initial tabula rasa proved as useful as had been Charles Darwin's similar protean quality as he boarded the Beagle.

Arriving at Rand in 1951, Allen Newell had been given considerable latitude in his choice of research topics. Newell himself put it more bluntly: "I was floating around doing whatever I damned pleased" (196). As we saw, this was a customary personnel practice (197). Newell used the time to get to know his intellectual surroundings. This does not seem to have been difficult to do: an engaging mind will be engaged by others as well. He was particularly intrigued by Cybernetics, and also learned something about game theory and worked for Oskar Morgenstern (198). He pondered going into this field, and returning to graduate school at Princeton, but ended up doing neither. However, he found his niche in psychology, specifically in organizational behavior, in the form of a dense study of human interaction with a continuous flow of information. Newell became a member of the American Psychological Association in 1952. He worked as a psychologist for the rest of his immensely productive life.

Rand allowed the option of small-scale independent work. Newell began by trying to unravel the collective decision-making process of a small group, but was only able to get mathematicians as test subjects. This did not work at all. Instead of solving the standardized-test-type problems collectively, they invariably solved the problems by themselves. This is not what Newell wanted, and he determined that he needed environments that were much more complex. Fortunately, the SRL Project, funded by the U.S. Air Force, was getting underway- and this would be an opportunity to meet Herbert Simon. This meeting, in turn would provide the best opportunity for optimizing, rather than satisficing, the SRL project's insights:

"...simple tasks were not the right environment for studying organizational behavior- you had to make the task environment

much richer, much more realistic, and you'd get genuine psychological behavior only out of environments that were too rich to allow the thinking human to think his way through and understand all the possibilities. And so we went from the little itty bitty one to the forty-man organization with a total simulated input, the air defense direction center." (199)

Herbert Simon had also been drawn into the SRL, albeit from a much more lofty perspective. He was brought in to examine the organizational problem of the SRL, and met Newell on his first visit, in February 1952 (200). The SRL changed his view of computers:

" I was familiar with computers- I'd wired some boards in my time, and I'd given lectures to businessmen on the implications of computers for business. But that Air Defense Lab was really an eye opener. They had this marvelous device there for simulating maps on old tabulating machines. Here you were using this thing not to print out statistics but to print out a picture, which the map was. Suddenly it was obvious that you didn't have to be limited to computing numbers. You could compute the position you wanted, a spot to appear on a piece of paper. You could print pictures, with things that weren't even a modern computer, just old card calculators." (201)

Simon returned on many occasions to meet with colleagues, and later, to keep abreast of the SRL and talk with Newell. They habitually drove together to the Air Force Base to see the SRL project in progress. Over the course of many conversations they reached "the idea that the computer could provide the formalism we were seeking- that we could use the computer to simulate all sorts of information processes and use computer languages as formal descriptions of those processes" (202). This seems to have been particularly cemented late in the summer of 1954, when Newell and Simon went out to observe a three-day exercise at March AFB (203).

" Our first conversation starting out from the parking lot was about the interpreter in the 701. But further on I can remember us saying well if we were really going to have a good theory of what goes on in human problem-solving, why not simulate it on the computer ?" (204)

They began to schematize problem-solving techniques by studying chess-playing, a domain that they would return to for at least the next decade. At about the same time, as we saw in Chapter Six, Newell and Shaw began to discuss designing the chess machine. This led to the IPL list processing language, and to larger discussions concerning human problem-solving (205).

Simon indicates that the perception of a need to mechanize thought was generally been a gradual process, rather than a lightning bolt (206). The gradual nature of this epiphany would seem to follow from the years of the early 1950s during which Simon became concerned with problem-solving. But Newell referred to a swift revelation of the possibility of computational complex information processing.

4. “Newell, Shaw, and Simon”

Working with Simon on the schematization of thought processing, Newell began to think about this topic in a more theoretical manner. As one of the designers of the SRL simulation, Newell began to work on computer programming in a more practical manner. In this capacity, he worked closely with Clifford (J.C., or ‘Cliff’) Shaw, the assembly programmer for the Johnniac (207). Newell referred to Shaw as “a non-talking person” (208). This is odd, given that Shaw talks a good deal in interviews such as those conducted by the Smithsonian in 1990 (209). No matter laconic or not, Shaw became a close colleague and the major programmer for Newell and Simon’s early cognitive science and AI work for the next decade, and often referred to together as Newell, Shaw, and Simon- hence ‘NSS’. The lack of information-processing view of the phenomenon being observed was a stimulant to Newell and Simon’s ongoing study of this capacity (210). Newell and Shaw together developed clusters of radar screens, which SRL used to simulate an actual Soviet attack. Newell, Shaw, and others arranged for the radar screens to depict a putative Soviet attack as realistically as possible:

“ We got all the flight plans from up there. We modeled the actual flights. We had the frequencies right. We worked on this map of the area. We got as much as possible of the actual air situation into the simulation.” (211)

Like Newell, Shaw, and Simon, the SRL project carried on well beyond its initial fruit. It produced several things: prosaic technical data on how to design control system displays, and

abstrusely theoretical responses such as Newell and Shaw's work in AI, and even more money for Rand. SRL led to an Air Force contract for the training of air defense specialists. By the mid-1950s, hundreds of people worked for Rand on this contract, and the SRL had turned into Rand's Systems Development Division. In turn, in 1957 this division was spun off to become the SDC Corporation. SDC operated the Dew air defense system and was a major Federal military systems contractor throughout the Cold War (212).

5. Oliver Selfridge, Signal Booster

“ One of the things I should emphasize about my life is how incredibly lucky I've been in meeting people. As an undergraduate just before entering V12, I roomed with Walter Pitts and Jerry Lettvin, which led me to Norbert and to Warren McCulloch... At that time, Warren was doing experiments in the University of Illinois labs in western Chicago, mostly on cats, checking neuronal connections etc.” (213)

Rand, as we saw, was significant in drawing in a large cast of significant intellectuals. Its intermittent visitors as well as its staff were often significant in its breakthroughs. One such meeting was that of Oliver Selfridge and Allen Newell.

Newell attributed his “conversion”, mentioned earlier, to both his conversations with Simon, and to his association with Oliver G. Selfridge (1926-2008), the Lincoln Laboratories research scientist who was mentioned in Chapter Three. His position as a founder of AI is equivocal. However, he greatly influenced Newell, who considered him an unsung cofounder.

Selfridge's training and pursuits resembles that of the Cyberneticists, although it is clear that he thought widely across disciplines. His father, an American department store executive, moved to England and founded Selfridge's department store. Selfridge's father was born in Britain, but left for the United States the day after the Blitz started (August 1940). He studied at MIT during the war years, graduating just as he turned nineteen, and enlisted in the Navy at the best possible time- that is, immediately after World War II's hostilities ceased. After two years as an electrical engineer at the Signal Lab Corps in Monmouth, New Jersey, he became the leader of the Communication Techniques group at

Lincoln Labs (214). At MIT, he worked for Norbert Wiener; as Wiener's assistant in 1948, Selfridge checked the page proofs of Cybernetics (215). He also became close friends with neurologist Walter Pitts (216). Late in the 1940s, McCulloch and his camp moved to MIT, where they were resident at the RLE for the next two decades or so. Selfridge was drawn into the "extended yet intimate scientific family of Rosenblueth, McCulloch, Wiener, [and] Pitts," although he was not invited to any of the Macy conferences (217).

Through these affiliations, Selfridge was more closely exposed to neurology, and to physiology, than were most others present at AI's creation. But while Wiener, McCulloch, Pitts, and the others were demonstrably uninterested in a cognitive approach to the mind, Selfridge appears to have had a broadly ecumenical interest in intelligence in many forms. He published actively in the field of pattern recognition throughout the 1960s and continued to collaborate with and advise early AI research at MIT as well. Selfridge invented the Pandemonium program, which devised subprograms that acted as agents (the first implementation of the concept), and was an advisor to Thomas Evans' "heuristic program to solve geometric analogy problems" (1964), an early MIT AI dissertation (218). He did some similar problem solving programs himself.

Selfridge, at the time a project manager at the Lincoln Laboratories, visited Rand in mid-November 1954, to lecture on the MTC, or Memory Test Computer, on which he was trying to emulate pattern recognition. Selfridge, and others at the Lincoln Laboratories, was working on a program that "learned" in the field of pattern recognition. Its domain was mechanical reading of letters and simple geometrical figures, which it would recognize by reading through a photocell and comparing the given samples to extant statistical norms. This was not 'learning' as AI, or at least the early NSS CIP learning, defined it. However, the pattern recognition research indicated that machines were capable of carrying out complex 'mental' tasks, based on very rudimentary steps.

Selfridge's 1954 lecture at the Rand Corporation seems to have been the catalyst for Allen Newell's epiphany of the future of symbol processing. Conversations with Selfridge and this lecture

led to what Newell described as a 'conversion experience'. His confusion and indecision were suddenly replaced by zealous clarity. In the hour following that lecture, Newell sought out Clifford Shaw and explained the concept of symbolic processing by computers to him in a frenetic hour-long monologue: "conversion experiences are all the same" (219).

Was Selfridge AI's Fifth Beatle? Not quite, since like the Beatles' own early manager Brian Epstein he dropped out of the band so early. Selfridge withdrew from active leadership in the AI community by the mid-1960s and worked at the Lincoln Laboratories through the rest of his life (220). Perhaps we might think of him, in Cybernetic terms, as a signal booster.

5. Newell's Definitive Turn and the Logic Theorist

Changed by the watershed intellectual encounters with Simon and Selfridge, Newell turned toward complex information processing- the psychological shade of AI. The end of 1954 was what Newell later called "the open point" in his life (221). Several paths welcomed him. Among the obvious ones, given what we know about him at this moment, all would guarantee steady employment and material success. The straight and narrow option was to return to graduate school at Princeton, study game theory further, and join the intellectual cadre of the Cold Warriors, at Rand or at a major university. Another option was pursuing Cybernetics, building automata which took their mechanical bases from neurology and physiology. Newell was indeed drawn to the McCulloch-Pitts thesis, as indicated by his intellectual companionship with Selfridge. However, his interest in information processing was more of the cognitive than of the wet-brain sort. Newell had also been offered a position as a junior fellow at the newly formed Center for Advanced Studies in the Behavioral Sciences, in the hills to the west of the Stanford campus. The final choice and most risky choice was pursuing the as yet nonexistent field of Complex Information Processing, taking a significant salary cut at a university far away.

This path was the one which Newell took. At the end of 1954, he wrote the chess paper discussed in Chapter Six, and he started to work systematically on heuristics as a means of problem-solving. He arranged to move himself and his wife and child to

Pittsburgh to join Herbert Simon at the Carnegie Institute of Technology in the Spring of 1955, although he spent several more summers working for Rand in Santa Monica. With the help of Rand manager Paul Armer, Newell tele-commuted to Rand for another six years. Even at this early date, telephone lines could carry digital signals. With specialized equipment, Newell and Shaw collaborated remotely for years on their next projects, the Logic Theorist, the chess playing program, and later, the General Problem Solver (222).

6. The Concept of Heuristics and the Logic Theorist

“ Inventor’s paradox: The more ambitious plan may have more chances of success.

This sounds paradoxical. Yet when passing from one problem to another, we may often observe that the new more ambitious problem is easier to handle than the original problem. More questions may be easier to answer than just one question. The more ambitious plan may have more chances of success provided it is not based on mere pretension but on some vision of the things beyond those immediately present.” (223)

Newell, Shaw, and Simon had come up with excellent, protean ideas. But how could they be put into some useful form ? Newell and Simon needed an object of study. Newell and Shaw were already discussing computer languages, which they would subsequently write themselves, as their vehicle. But the appropriate object for testing out the idea of heuristic problem solving by computer was not entirely clear. They began to pursue the proper problem.

Clearly, they would need to implement the meat of these conversations in computer programs. But how ? By 1955, NSS were trying to figure out rudimentary tasks that the computer could be given to do. They went over several possibilities, including geometrical theorem-proving, chess, and logic. When Newell returned to Pittsburgh in the Fall of 1955, he was determined to program a computer to play chess. But he and Simon began to consider other prospective objects of applications as well, and began to meet each Saturday to discuss them. In October of 1955, Newell and Simon were attending an Institute of Management Sciences meeting in New York City. While strolling in

Morningside Heights, adjacent to Columbia University, Simon suddenly realized that solving proofs in geometry would be a better testbed than solving chess-playing as a problem- his equivalent of a sudden epiphany. He and Newell thus turned to geometric proofs (224).

Once they started working on geometry, they found diagram depiction by computer language to be very difficult. Eventually they reached symbolic logic as a first petri dish, perhaps in part because Simon had a copy of Whitehead and Russell's Principia Mathematica (1935) in his office, and was familiar with Carnap's formalism from his days at the University of Illinois at Chicago (225). They used this highly restricted formal language, which could be rendered into Boolean and binary form, and started to work in pseudocode.

Another concept that was clearly quite central was that of heuristics, the idea of problem-solving through rules of thumb, or human-like or rhetorical means. This was one of the central heuristic concepts in *How to solve it: A New Aspect of Mathematical Method*, by George Polya, the mathematician who had inspired Newell in his college days at Stanford.

One final immediate intellectual influence bears mention. Examination of the work of De Groot led Newell and Simon to begin to use protocols, the talking evidence of problem-solving processes. These have proved important for both psychology and AI. This was clearly a cocktail with high alcohol content, and they were more than ready to drink it.

The Logic Theorist consisted, then, of an effort to implement on a computer some of the proofs that had been carried out by Russell and Whitehead in the Principia. The various logical statements would correspond to machine states. To make the machine 'think' in this way requires, first, the proof itself. These were usefully canonically enshrined in the Principia, along with the fundamental axioms of manipulation of logical symbols.

Thus LT engaged in a kind of algebraic problem-solving without any actual numbers. It was aided in this in that it used a sort of racing-car logical symbol system, the sentential logic, which lacked propositions and predicates about specific substantive things (e.g., 'Fred is a frog'). Current AI uses the first order predicate logic, a far more extensive logic; but the more limited

system was enough to get started. NSS created the first of several successive computer languages- called 'IPL' (information processing language). Starting with a theorem, LT proves it by engaging in successive [logically] legal operations, usually operations of substitution, until the theorem's expressions are reiterated. There are other logically legal operations, notably replacement and detachment, in the Principia axioms, but the preponderant operation carried out appears to be substitution that reduces the search space (226).

The computing space for this task was mercifully small, although adequate to be a difficult task for a human and to require solution times of minutes rather than seconds. Prior to the program's implementation on the JOHNNIAC, the problems were simulated with index cards, with cards representing the program (the subroutines) and the memory (the axioms of logic) as provided in the Principia (227). Blind substitution would eventually lead to proofs, particularly if time was cheap, regardless of the fact that it had no 'intelligence'. But NSS did better than this, and determined that their program would cache its solutions as well. This was significant because it provided a first step toward learning from one's mistakes and successes, rather than slipping on the banana peel in the same location repeatedly. The heuristics that LT implemented beyond simply using the stipulated logical operations include the caching of difficult subproblems. This alone dramatically streamlined the solution, making the heuristic method greatly more efficient than the blind British Museum Algorithm.

The value of such objects is clearly in the eye of the beholder. In his autobiography, Herbert Simon recounts that the Journal of Symbolic Logic rejected the computer-generated proofs as a publication on the grounds that these proofs were already well-known. Thus, the Logic Theorist was not at all novel ! Bertrand Russell, however, was flattered and impressed, and wrote to Simon to say as much (228).

Computational problem-solving, whether it emulated humans or was simply more 'rational' in the operations research and flow-charts sort of manner, was the wave of the future. But the Logic Theorist was running on the computer technology of the present, and this presented an ongoing struggle. Considering just how

early it was in the history of digital computing, it is a bit amazing that NSS thought of the Logic Theorist and the other programs at all. It would seem that the SAGE and SRL projects had afforded this look at things to come. SAGE had offered its own glimpse into computing in the future, with the possibility of far more interactivity than was available in the mid-1950s. The SAGE or SRL user gazed at CRT screens rather than at machine registers or printed cards. Moreover, SAGE users were able to use direct input, that is the typewriter keyboard, rather than punch-cards. Output, at least in the case of the LT work on the Johnniac, was in the form of a 40-column numeric printer, cumbersome to decode. An alphanumeric printer was not available until 1957 (229).

SAGE, in the form of the extravagant SRL project, had suggested that programming could be made much more accessible, and applications made much more elaborate as well. But now NSS were no longer working on the SRL project, but in the real world- or at least at Rand. The SRL-type control rooms, with the most modern computing available on planet Earth, was now found only at the dozen-odd secret SAGE installations and at the SAGE research outposts. NSS, perhaps thinking in terms of computational possibilities which would not be realized more widely for at least a few more years, had to content themselves with the Rand Johnniac. Having the Johnniac to confront was in itself an indication of vast privilege; Johnniac was one of the very first digital computers. But this did not mean that it was easy to use:

“ Before the Institute [Institute for Advanced Studies] computer was finished, Willis Ware left to become head of a group at the Rand Corporation... built the Johnniac... which passed its acceptance test in March 1954. Meanwhile, the AEC’s laboratories, Los Alamos, Argonne, and Oak Ridge also built copies... The RCA group under Zworykin and Rajchman opted to build a storage tube that became known as the Selectron, and which was in fact used in the Rand Johnniac.” (230)

Newell, Shaw and Simon’s team worked neatly during this time period. Its internal division of labor was clearly stated and effective. Simon was typically responsible for the strategic organization of the early computer programs, while Newell and Shaw developed the computer languages (231). Simon did,

however, learn to program the IBM 701 in the summer of 1954, and distributed its manuals to his students (232). Shaw was the programmer, although Simon rarely spoke directly with him. Newell acted as the middleman (233). However, Simon never fails to acknowledge Shaw's contribution. The team came up with a hand-simulated program in mid-December, 1955.

The programming necessary to enable the Johnniac to run the Logic Theorist was not completed until August 9, 1956. The hardware was not trivial, but was secondary to the vision of the computing. "Of course LT wasn't running on the computer yet, but we knew precisely how to write the program" (234).

Meanwhile, Simon was diligently proceeding with hand simulation with index cards, and Newell and Shaw were doing their work by computer. Decades before telecommuting was commonplace or even possible for most people, Newell and Shaw collaborated through an early teletype (an electronic communications device). They could not have done this but for the monies spent by Rand; the phone bills reached eight hundred dollars per month. Fortunately, money was rarely an object as far as the Rand Corporation was concerned.

In his CBI Oral History interview, Newell relates the great noise made by the model 28 Teletype in his Pittsburgh apartment, communicating through telephone wires with Cliff Shaw in Santa Monica many hours per day:

"It was god awful, terrible clanking. Everyone around the area thought we were the bookie place. They really did."

Simon pinpointed December 15, 1955 as "the birthday of heuristic problem solving by computer", the day on which he finally got his index cards properly lined up (235). NSS completed the first 'paper' simulation of the Logic Theorist during the Christmas 1955 vacation, using paper and human beings:

" While awaiting completion of the computer implementation of LT, AI and I wrote out the rules for the components of the programs (subroutines) in English on index cards, and also made up cards for the contents of the memories (the axioms of logic). At the GSIA building on a dark winter evening in January 1956, we assembled my wife and three children together with some graduate students. To each member of the group, we gave one of the cards, so that each person became in effect a component of

the LT computer program- a subroutine that performed some special function or a component of its memory. It was the task of each participant to execute his or her subroutine, or to provide the contents of his or her memory whenever called by the routine at the next level above that was then in control.

So we were able to simulate the behavior of LT with a computer constructed of human components. Here was nature imitating art imitating nature...Our children were then nine, eleven, and thirteen. The occasion remains vivid in their memories.” (236)

At the commencement of classes in January 1956, Simon announced to his students, among them the young Edward Feigenbaum, that, “Over Christmas, Allen Newell and I invented a thinking machine” (237). And even before the LT had appeared on a computer program, Newell and Simon were publishing its achievement. The LT simulation was first published in Rand Report 850 on May 1, 1956, and was first presented by Newell, at a seminar called “Current Developments in Information Processing”, the next day in Washington, D.C. This defines the first publication in AI.

Chapter 9. The Declaration of AI at Dartmouth

1. Introduction

AI appeared as a field in the mid-1950s. From far away, such events appear to have sharp edges. AI is commonly thought to have come into being at the Dartmouth Summer Research Project on Artificial Intelligence of 1956, usually called the Dartmouth Conference. At close range, we can see that it cohered throughout the early 1950s, with Minsky and McCarthy’s approaches to computing as a way to build automata and Newell, Shaw, and Simon’s Logic Theorist. This is especially the case when we consider the role of Cybernetics as a “halfway house” between Behaviorism and the Cognitive Revolution.

Yet 1956 remained the modal year for such events. This year witnessed the two conferences important for the naming and articulation of AI as the science and design of cognitive artifacts. The first, the Dartmouth Conference, has eclipsed in renown the

Symposium on Information Theory, held by the Institute of Radio Engineers at MIT, in September of the same year.

2. Organizing the Dartmouth Conference

Simply naming something as if this ensures its existence is grandiose. Yet in certain cases such gestures are more than politics or public relations; they may become a self-fulfilling prophecy, which helps to give the field greater credence and substance.

This was the case with the Dartmouth Conference of 1956, which brought together the four founders of AI for the first time. John McCarthy, at this time an assistant professor of mathematics at Dartmouth, was the chief organizer. McCarthy established that the topic would be Artificial Intelligence, defined as “the belief that every aspect of learning or any other feature of intelligence could be simulated” (238).

The organizing committee consisted of Minsky, a Harvard Fellow at the time; Claude Shannon, for whom Minsky and McCarthy had worked in the summer of 1952 at the Bell Laboratories; Nathaniel Rochester, manager of information research at the IBM Laboratories in Poughkeepsie; and McCarthy himself. The group’s joint proposal suggested that:

" A two-month, ten-man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth in Hanover, N.H. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it".

The Rockefeller Foundation gave the group the seventy-five hundred dollars in funds they asked for to cover rent, salaries, railway fare, and other expenses (239). The Office of Naval Research also provided some funds, although Martin Denicoff, the ONR's typically enthusiastic program officer, was uncertain:

" [The Conference was] so novel at the time that no one really knew what would come of it. In fact, I questioned that the participants had a good insight into what was going to happen..." (240)

Denicoff’s skepticism was well-founded with regard to any expectations of consensus among the conference participants.

Yet he proved to be misguided in doubting the eventual fertility of the meeting. By drawing together individuals with vastly different views and methods for approaching roughly the same topic, the organizers assured lack of harmony, but not lack of progress.

The topics of AI as originally envisioned were open-ended, and this clearly was its great strength. McCarthy's own proposal for research suggests training on simple tasks by means of 'trial and error'. This purportedly will result in the machine's construction of an abstract model of the environment. It is not clear how the former will generate the latter, but it is clear that 'goal-seeking' behavior is sought. The initial inventory of topics which the conference organizers wished to see addressed includes 'randomness and creativity', as well as 'abstractions', but creativity is described as 'controlled randomness in otherwise orderly thinking'.

However, the way in which this would be pursued was not clear at this time. In retrospect, particularly when one looks at AI in the 1960s, the predominance of AI as 'artificial thinking' during the third quarter of the century appears natural. But this belief is in itself an artifact. The majority of the metaphors for computation in the proposal are neurological and of a Cybernetic cast. The final list of attendees and invitees is something of a 'Cybernetic cast' itself (240). The list of invitees shows a very broad grouping of people from every relevant area: computer designers and programmers (John Backus, IBM; Robert Fano, the RLE; David Hagelbarger, Bell Telephone Laboratories; Nathaniel Rochester; Julian Bigelow, IAS); Cyberneticists of the mechanical sort (Ross Ashby, Donald MacKay); and of the neurological bent (Warren McCulloch and Walter Rosenblith from the RLE, and Oliver Selfridge) (242). There were also mathematicians (John Nash, from the IAS and Norman Shapiro from Rand) and electrical engineers (Claude Shannon), and one of the world's few cognitive psychologists (George Miller, Harvard). Finally, the first AI and computer-game playing people were invited- Newell and Simon, Samuel and Bernstein. Clifford Shaw- brilliant but self-taught and lacking degrees- was not invited.

Denicoff rightly saw AI as a novel advancement of ideas about computing. But there is no denying that it was an inside job, rather than a movement of the intellectually disenfranchised- the

guest list to the Conference confirms it. Even if their ideas were staggering, the founders of AI were extremely well-connected. Illustrating this merely takes some name-dropping. John McCarthy worked early in his career with John Kemeny, a frequent Rand consultant, mathematician and computer scientist and later President of Dartmouth University. Marvin Minsky's numerous affiliations with eminent and powerful scientists began with conversations with B.F. Skinner as an undergraduate, and include interest and support for his projects from the top leaders at MIT and Princeton. Herbert Simon worked with Richard Cyert, an influential economist and later President of Carnegie-Mellon. AI was a cutting-edge palace coup rather than a group of revolutionaries storming the gates.

A weighted average of the interests of the attendees would have steered the ultimate, or even temporary definition of AI, toward Cybernetics as it stood in 1956. But published words are more lasting than talk, and surely this is why NSS proved so influential in the longer run. There were many luminaries present, but there were not many papers. Only Newell and Simon and Trenchard Moore of MIT presented any finished work at all. Moore of MIT brought a program for theorem proving (243). Of the works presented, the research most advanced by far was the Logic Theorist. However, two other hand-simulation programs inadvertently resembled, in much less developed form, the logical theorem-proving of the Carnegie Institute professors.

Marvin Minsky brought with him mimeographs of his Euclid notes (244). Earlier in 1956, Minsky re-read Euclid's Elements and found that the hundreds of specific theorems could be distilled into a far smaller number of genres. Simulating theorem-proofs through pseudo-code on paper, he succeeded in deriving a computer compatible proof for one such theorem (in an isosceles triangle, AB, BC, and CA are all equal). While Minsky did not later pursue this work, his student Herbert Gelernter did. Minsky told Pamela McCorduck that he "considered the idea of heuristic search obvious and natural" (245).

3. Initial Implications of the Dartmouth Conference

A glance at the scenario presented by the conference will suggest that it would, like the first act of a play, immediately

evidence certain tensions. After planting their flag on the new shore so firmly, Newell and Simon wanted recognition. They already had an agenda as well as a computer program, of course, and publications and were bringing in their own graduate students, and an established academic program at their university and at Rand. They did not receive as much recognition as they wished (246). Neither Newell nor Simon, after attending for only one week, thought of the conference as a decisive event, because "Herb and I were totally consumed with our own path" (247). Simon asserted that he and his family had a pleasant vacation in New Hampshire (248), but that nothing was resolved or changed because of the event. John McCarthy wished to provide more than the title for the new field- however, the substance in terms of programs was provided by NSS. Since 1956, his contributions have been legion, but at the time, all that he brought with him, or took away, was the title. As McCorduck puts it,

" Neither Minsky nor anyone else had been able to extend beyond the trivial the neural model of human cognition promoted by McCulloch and his followers. McCarthy's hope of inventing a formalism to describe human thought, a calculus ratiocinator was looking more and more impossible...and then someone else got the prize." (249)

Minsky said that NSS seemed to be more interested in psychology than in AI. This situation was certainly exacerbated by Simon's characteristic reticence and reluctance to publicize himself. This was bound to lead to many such quandaries, and this was the first expression of intellectual and personal rifts to come.

Moreover, some of the problem may have been, as indicated earlier, in the design of the list of invitees. The Conference put together too many people, in too many fields, over too long a period of time. Broad disciplinary conferences work by bringing together many people over a longer period of time. But these events are often expository, intended for everyone to present their research, rather than intended to produce the outlines of a new field. This higher goal was what the Dartmouth Conference's original statement of intention indicated. So McCarthy was trying to do what the Macy Conferences had done, but was designing

the conference structure as if his planned event was an expository presentation of an existing field.

McCarthy has expressed disappointment that there was never a sustained, constructive intellectual encounter of the sort that he had wanted (250). Unfortunately, academic conferences often do not yield such encounters. The setting may be an excuse for presenting a paper and talking with established colleagues. This means that conferences can result in normal science, or further improvement of the ideas one is working on, with the people one is already working with- preaching to the incumbents, we might say. For instance, Newell and Simon conversed at length with John McCarthy concerning IPL, the early symbolic computer language (251). This was interesting, but Newell and McCarthy had had such talks a year before. By making the session so lengthy, rather than a one or two day-session, this seminar probably discouraged confrontation that just might have been intellectually challenging enough to seem appropriate for the parties present.

The presence of Walter Pitts, Herbert Simon, and Claude Shannon in the same room did not mean that they discussed their ideas at a high metaphysical level. In the case of Newell and Pitts, for instance, one can see that they had instead talked right past each other. Or at least, Pitts had in the recent past talked right past Newell. At the 1955 Western Joint Computer Conference, immediately after the Session on Learning Machines, at which Newell presented his paper on the problem of Chess, Pitts responded:

“ The speakers this morning are all imitators in the sense that the poet in Aristotle ‘imitates’ life. But, whereas Messrs. Farley, Clark, Selfridge and Dinneen are imitating the nervous system, Mr. Newell prefers to imitate the hierarchy of final causes traditionally called the mind. It will come to the same thing in the end, no doubt, but for the moment I can only leave him to Mr. Miller and confine my detailed remarks to others.” (252)

Pitts then launches into a lengthy, technical monologue on different neurons in the brain.

4. The IRE Conference

The Symposium on Information Theory at MIT, in the beginning of September 1956, was a fine coda to Dartmouth. Organized by

Claude Shannon, the Institute of Radio Engineers program included presentations of Miller's 'Magic Number Seven' paper, a paper on the transformational grammar by Chomsky, Newell and Simon with the 'Logic Theorist', and a seminar by McCarthy, in addition to appearances from most of the other major figures from the earlier Dartmouth gathering. Miller's paper established the apparent primacy of seven digits as the number beyond which human short-term memory typically erodes in accuracy (253). Chomsky's paper was an early work in his project to describe innate human linguistic commonalities (254). In all, it presented a good cross-section of the work in cognitive psychology, linguistics, and information theory of the time. This was made more comprehensive by the presence of Newell and Simon.

The issue of how to report on the Dartmouth Conference illustrates both McCarthy's determination to develop the field in his way (not that it was clear what that would be), and the insistence of NSS that their accomplishments be recognized as well. It was suggested that McCarthy report alone on the Dartmouth conference; Newell and Simon wanted a presentation as well. Walter Rosenblith, chair of the session (255), was clearly stuck between a rock and a hard place. No one on earth would wish to be caught between such mighty intellects and wills. His diplomatic solution consisted of one presentation by Newell and Simon, and one by McCarthy (256).

5. The Significance of the Conferences

The consequences of the Dartmouth Conference are significant despite its lack of resolution. It seems to have fortified McCarthy's already considerable resolve; it lent a name to the work of NSS, whether or not it was the name they had given themselves. The name 'Artificial Intelligence', coined by McCarthy, prevailed after the Conference, regardless of the existing work by Newell Shaw and Simon in Complex Information Processing. Did this name choice really matter, especially given the emphatic dedication of the scientists to their chosen paths? It appears that there's less to the name than one might think.

The title 'AI' itself was more ambitious than wise, though. Turing used the term "thinking machines". Marvin Minsky has pointed out that 'we don't usually name fields for their

aspirations', but for their subject matter or their function. (Apparently, no one paid attention to this advice when A-Life was named, either). For instance, Genetics, 'the biological study of heredity' is not called 'people-and-animals-by-design'. Moreover, a title that refers to the scientific as well as to the engineering content of a field is certainly wiser. It does not implicitly indicate that the field in question is the technical auxiliary to more universal questions. Simon's nomenclature seems more strategically savvy and more semantically appropriate. It invests less bravado in its epistemological claims. CIP would certainly have been a more fortunate choice of names, although the substance has doubtless not differed.

However, the name of AI stuck. Its origins, strictly speaking, belong to several different individuals. The field, born as 'complex information processing', originated during the summer of 1954, in the course of the collaboration between Newell and Simon, not at the Dartmouth Conference in 1956. The answer is the same if we define AI as cognitive emulation. If we define AI as the computational formalization of logical reasoning- as opposed to thinking of the human variety- then McCarthy's published work on automata would seem to claim priority, with the Logic Theorist as nevertheless the first "AI" program. The first real discussion of 'machine intelligence', or 'thinking machines', was proposed by Alan Turing in the late 1940s, as was the abstract concept of software and high-level program languages. Moreover, should we define AI as computer languages, or the infrastructure to instantiate such formalizations, it is Newell, Shaw, and Simon once more. Under a looser definition of AI, and accepting the move once removed from cellular automata to artificial intelligence, a larger group of people, including Turing, Selfridge and Shaw and especially the four acknowledged founders, should share credit.

Simon readily acknowledged his concern with the recognition of priority in scientific discovery, specifically with his and Newell's discovery of AI (or complex information processing) prior to the Dartmouth Conference:

" Priority in science is important, as Newton said [sic]. This does not mean that he was a nice man- or that I am." (257)

Chapter 10. The Inexorable Path of Newell and Simon

Introduction: Constructing a Research Environment at the Graduate School of Industrial Administration

The form of complex information processing presented by Newell and Simon was the best-developed one of the mid-1950s. But the undisputed success of this model during AI's rapid development during the 1960s and 1970s may obscure its relative novelty circa 1956. At the Dartmouth Conference, most of the participants had been drawn from the Cybernetics and electrical engineering cohort at various MIT laboratories. When Newell and Simon returned to Pittsburgh, the CIP approach to AI was less favored. During the years following 1956, the energetic and steady pace of Newell and Simon's work bore fruit and firmly established their approach as part of the state of the art.

Newell and Simon had already built the outlines of their research project, which they put into practice through publication and presentation. The construction of a discipline entails the development of a descriptive vocabulary- and NSS did a great deal of vocabulary-building as part of discipline-building. (We will introduce some of their novel vocabulary in this chapter). Working with Cliff Shaw, but increasingly as a pair rather than a trio, Newell and Simon trained graduate students, taught, and worked to expand the infrastructure of the GSIA and eventually establish a computer science department. They procured research money- Simon pays tribute to Newell's unique talent for "knowing where to put the decimal point on proposals" for such monies (258). They presented their work at gatherings of computer scientists, of the Operations Research Association, and of psychologists. For the latter, they arranged special summer seminars at the Rand Corporation in 1958, 1962, and 1963. Using computers donated by IBM, they developed successively better computer languages. IPL went through at least five versions. (The latest version, OPS-5, remained in use for decades). The IBM computers would later be replaced by DEC's more user-friendly PDP (259). They published furiously, jointly and separately, and played an entire orchestra of disciplines at once. Simon continued to work on the theory of bounded rationality in economics, although his devotion to this

work was somewhat diminished by the field's lack of interest. Newell and Simon's work with psychologists appeared in the volume Representation and Meaning (260).

Simon's graduate students set to work at various projects. Among these were "building a program that used heuristic search to balance assembly lines (to find the best arrangement of workers tasks and work stations) while Geoffrey Clarkson constructed an expert system (as we would call it today) for choosing stock portfolios for bank trust accounts." (261)

Finally, graduate students turned into dissertation advisees, in some cases later longtime colleagues. Simon's cultivation of a young GSIA student named Edward Feigenbaum led to the Elementary Perceiver and Organizer (EPAM), designed to study how the brain chunks material in short-term memory. Feigenbaum, subsequently a Stanford professor, would later figure prominently in the history of AI as both a researcher and a policy advocate (262).

Herbert Simon Wearies of Economics' Inflexibility

Much as it is redundant to repeat it, Herbert Simon held the scalpel at the cutting edge of numerous fields in the Twentieth century. Simon's prodigious opus astounds; however, he did not work at the same pace and same intensity of achievement in every field for the duration of his work. His path led from city management, to political science, to economics, then into AI, in a trajectory of increasing abstraction as he considered the ways in which human beings made decisions. Finally, his work in AI considered human information processing embodied in the program itself.

As was discussed in the third chapter introducing the founders of AI and their generation, Simon was trained as a political scientist at the University of Chicago (263). During his years as a professor at U.C. Berkeley, then at the University of Illinois in Chicago, he moved into the realm of microeconomics by questioning the assumptions of individual rationality explicitly stated and universally assumed in that field (264).

It is not surprising that Simon gravitated to the intellectual high ground and focus of elite intellectual interest so rapidly. He was well aware that the most pre-eminent field in the social

sciences was theoretical economics. In 1944, John Von Neumann and Oskar Morgenstern published the “Theory of Games and Economic Behavior”, a classic work concerning caveats to rational decision-making which are imposed by the informational asymmetries of game-playing. McCorduck tells us:

“ When Herbert Simon saw an advertisement for the book... he felt a flush of envy so great he could remember it vividly thirty years later.” (265)

Such profound emotion evoked strong responses. This examination of empirical facts led inductively to subsequently more abstract fields, in an apparently cumulative process. Administrative Behavior, Simon’s critique of micro-economics’ assumption of universally rational man, appeared in 1947. Simon proposed that rationality in actual human economic behavior is quite limited, contrary to the assumptions of human rationality in classical microeconomics. The costs of search, or finding out all economic facts about commodity X , may be too high to allow people to get as much information as they would need to know about all prices. Thus, people “satisfice”: they get along with the limited information, rather than getting the best possible product or price. Except in (rare) perfect markets, consumers satisfice rather than make the optimal or perfect, decision. Rationality is bounded because of the structure of the informational environment- most of us don’t have time or energy to search for the lowest price (266). Thus, economic actors are “boundedly rational”. They make do with the information that they have, and are usually thrifty or wise in their business decisions only in this context. Simon later applied the concept of decision making to firms, purporting that such institutions are, like individuals, boundedly rational (267). The market-clearing assumptions of NeoClassical macroeconomics economics require that people engage in search costlessly, and that they buy everything at the lowest price (that is, markets “clear” as everything is sold). The caveat of boundedness was alternatively recognized as brilliant and relegated to footnotes. However, the influence of Administrative Behavior on economics was specifically mentioned when Simon was awarded the Nobel Prize.

Simon’s background in the social sciences was intellectually decisive in his work in AI, in ways not immediately evident. Unlike

McCarthy, Minsky, and Newell, Simon never studied math as a formal major or graduate field. This meant in turn that he learned the Cybernetic view of intelligence as biological from afar, as an adult, rather than from his professors as a much younger person. Thus Simon never saw cognition as a neurological event, which could best be modeled through local self-organizing units. His considerable knowledge of science and mathematics was gained on his own, and so he was not much influenced by the ambient culture of Cybernetics during the early Post-War years. Instead of thinking of human intelligence as a species of neural inputs and outputs, he began by looking at more substantive and complex thoughts- which were best incarnated through symbol processing. Problem-solving was thus the single topic on which the great brunt of Simon's intellect weighed. Even before he thought of using a digital computer as a problem-solving symbol processor, Simon saw problem-solving as directed search. The notes accompanying this paper indicates that the issue was "thought through" in 1952, and the text indicates the significance of search:

" In most global models of rational choice, all alternatives are evaluated before a choice is made. In actual human decision-making, alternatives are often examined sequentially. We may, or may not, know the mechanism that determines the order of procedure. When alternatives are examined sequentially, we may regard the first satisfactory alternative that is evaluated as such as the one actually selected." (268)

However, the pace and acceleration factor, so to speak, of the field of economics was not equal to that of other fields. Simon was, ironically, active in the highest circles of the economics profession in the United States in the Forties and Fifties. The Cowles Commission, for instance, was a distinguished committee devoted to lending more rigour to the profession as a whole. Its influence was largely to increase the importance of econometrics (269).

Simon was the man of the hour, in the right place at the right time, along with several others in the mid-1950s for the formation of AI (270). But in economics, he was so far ahead of his time that he was out of step with the profession as a whole. He was ready for this party, but no one else had showed up yet. Behaviorist

approaches which would require recognition of the less than perfect rationality of human behavior and knowledge were not popular for decades after Simon did his work. Simon would not win the Nobel prize in economics until 1978. Eventually, during the 1950s, he started spending less time on economics and more on Complex information processing, at the Rand Corporation and at Carnegie's GSIA (MOML). Economics' loss was to be AI's gain, as Newell and Simon turned to psychology.

The Initiation of Cognitive Psychology

As we saw in Chapters Two and Seven, problem solving had not been an ongoing topic of cognitive psychology in the years of AI's gestation and birth. Yet despite impediments the field of cognitive psychology appeared of its own accord as early as 1950, from the malcontents of Skinnerian psychology departments. This was an indigenous movement on the part of scholars who wished for better descriptive concepts for the empirical phenomena of cogitation. But it was also precipitated by foreign provocateurs, so to speak, such as Freudian personality theory, and by the failure of information theory to say anything about semantics (271).

The discontent was longstanding. Even in the 1940s, U.C. Berkeley's E. Tolman had suggested modifications to the existing crude stimulus response (SR) theory:

“ As Stimulus-Response theories came to be modified to take into account the subtle events that may occur between the IP of a physical stimulus and the emission of an observable response, the old image of the stimulus response bond began to dissolve, its place being taken by a mediation model. As Edward Tolman so felicitously put it some years ago, in place of a telephone switchboard connecting stimuli [sic] and responses it might be more profitable to think of a map room where stimuli were sorted out and arranged before every response occurred, and one might do well to have a closer look at these intervening 'cognitive maps'.” (272)

As with many other fields, innovations in empirical research began in Cambridge. The Harvard Center for Cognitive Studies, initiated in 1951 as the Cognition Project in the Laboratory of Social Relations, was nurtured by Rockefeller Foundation and Ford

Foundation grants and time at the Institute for Advanced Studies (273). The Center eventually attained hard-money status. In 1956, Jerome Bruner, Jacqueline Goodnow, and George Austin published *A Study of Thinking*, an academic work addressing human usage of categorization. Another major development was the appearance of Noam Chomsky and the subsequent development of linguistics. Chomsky emerged with a vengeance from the University of Pennsylvania about 1954, spent several years as a Harvard Fellow, and published *Syntactic Structures*, the first and resounding statement of his theory of innate human linguistic tendencies. Chomsky insisted in his work that restricted and simplistic behaviorist models of linguistic processing of language could not account for the complexity of actual human grammars. Indefatigable then as now, he commenced a sustained attack upon B.F. Skinner in the-mid 1950s.

This was virtually the only research being done in clinical psychology. However, two bodies of work at the time eventually influenced AI to various degrees. In Europe, the Genevan Jean Piaget was beginning his lifelong study of the way that young children construct categories and concepts- “genetic epistemology” as he called it (274). The term clearly translates badly- it refers to the genesis, or development of epistemology, rather than to genetics per se. Piaget stressed the development of categorization, and insisted that it could and should be studied in great detail. Thus this work was in essence and assumptions intrinsically opposed to Behaviorism. The same is true of the careful empirical work of the Dutch de Groot. Adrianus Dingeman de Groot in the Netherlands wrote the classic *Het denken van den schaker* (1946) translated as *Thought and choice in chess* (1965), which closely examined the strategies of chess players. The work is essentially an examination of protocols, although de Groot did not use that term. Protocols would figure heavily in Allen Newell’s observations of the sequences of items in problem-solving, and protocols as the concept of canonical forms of problem-solving would become a standard term in computer science. Herbert Simon indicates that his own most frequent references among American psychologists in *Administrative Behavior* (1947) were to William James and Edward Tolman (275). However, subsequent to the formation of cognitive psychology, he and Newell became

part of the field. Their summer lecture sessions for psychologists at Rand in 1958 and publication in *Psychological Review* were only part of their sustained engagement with this field (276).

In the thirty years before this, psychology had been fixed in the physiological rather than the psychological. The appearance of close, fine empirical academic studies which concerned the nature of concept attainment, categorization, how many things people could recall at one time, how the order of these things mattered, how they solved problems and how they broke down sentences into meaningful bits, began to greatly alter the entire discipline. AI's earliest practitioners would both draw from and contributed to the study of thinking, which came to be known as the Cognitive Revolution.

Toy Problems and the General Problem Solver

Simon calls chess the drosophila [fruit fly] of AI (277), but certainly the problems in GPS served the same function. Newell and Simon were well aware of the Logic Theorist's deficiencies, particularly those pertaining to cognitive verisimilitude. The Logic Theorist was succeeded by the ten-year General Problem Solver (GPS) project. In their next opus, they explored and expanded all of these concepts. They devoted themselves again to the dual challenges of cognitive science as such and cognitive emulation through technology (278). The project, still lacking a name, was first presented at a 1958 symposium, and the same year at a Rand summer seminar. GPS was not published until 1962 (279). In this project, Newell and Simon succeeded in creating most of the basic concepts of Classical Physical Symbol System Artificial Intelligence: search in problem-solving, including goal states, initial states, state-space evaluation, operators with self-consciousness at efficiency, methods of constraining search, learning through records of search history, and the idea of attaining more efficient- that is, cheaper- search through means-ends analysis (280).

Practicing with toy problem spaces, the General Problem Solver introduced a 'bigger' and more difficult search. It relied on iterative usage of both human protocols and computation (281). The tenacity with which Newell and Simon pursued the idea of emulating human intelligence was rewarded by the number and

usefulness of these working tools. They attempted to develop programming heuristics (word used loosely) which were adequately generic to be applicable regardless of the task environment. The General Problem Solver took on general and schematic search spaces, such as those in games somewhat simpler than those found on the Scholastic Aptitude Test. These problems were still very useful in teasing out the issue of the protocols involved in solving a problem. Somewhat derided when demonstrated because of their schematic nature, the games constituted the bread and butter of AI for the decade between 1956 and 1965.

In the 'Tower of Hanoi' (282), blocks of different dimensions must be moved across a row of several pegs in order to reach a particular size and proximity order. In 'Monkey and Bananas', a monkey must figure out how to reach a banana suspended far above its head. The solution is to use the empty box in the corner as a step-stool. In 'Cannibals and Missionaries', the program as protagonist is faced with the dilemma of transporting three of each of the latter across a river in a canoe that holds three passengers. Presuming that the cannibals are hungry and the missionaries unarmed, the program must never leave a cannibal and a missionary alone together on the home side of the riverbank (283).

These are useful problems because they strip down the factual specifics to ones which can be described in a few sentences. They can be easily reduced to digital form and expressed in the terms of the basic Boolean operators. (That is, 'and', 'or', 'not', and 'exclusive or'). Problems that lend themselves directly and simply to traditional forms of symbolic knowledge representation; that is, the form of representation is selected in advance. This was appropriate for a first approximation of learning how to solve problems. Once the concepts of blind or domain-free search had been better developed, toy problems would need to be supplanted by problems involving more substantive domains. However, in the early versions the structure of searching through a problem space itself needed solidification.

The basic components of computational search were articulated through the GPS research during these years. The program is presented with an initial state, from which the

objective stipulated is to reach a goal state. The possible variables make up a toolbox or options which the program may take to reach this goal. The program uses these operators to move through a search space corresponding to different possible paths toward a solution. Just as in economics or everyday life, the objective of the program is to minimize the time-computational in this instance- spent in successfully executing the search. The search proceeds through the application of the operators, as the operators are searched through a graph (or tree), and evaluates at each successively deeper node of the branches thereof whether this move brings the goal state any closer.

NSS had thought earlier of efficiency as a desideratum of efficiency in simulated cogitation, but had only begun to instantiate it in the mechanics of the Logic Theorist. The introduction of means-ends analysis was a critical innovation in AI. The operators in this means-ends metric architecture included preconditions without which they cannot be fired; transformation function; and the assessment of the differences reduced (284). The 'self-conscious' operators implemented in GPS include an evaluation function to assess, after their implementation, whether they have advanced the path's proximity to the goal state. This function is also called the matching function, in that it matches the current state against the goal state. Not only is this 'penny-pinching' feature basic to efficiency in search and learning, it demonstrates the persistent grasping at the generic features of intelligence on the part of Newell and Simon. Philosophically inclined, they were aware of the intelligent behavior as a constant evaluation of the entity's (the program's) status relative to the environment.

Cybernetics had defined intelligence as goal-oriented behavior. Newell and Simon redefined intelligence as problem-solving for fulfillment of a designated goal. Thus they transposed the earlier generic and mechanical definition for one involving cogitation. More precisely, they used the concept of feedback, or iterative adaptation to one's environment, as defined by the Cyberneticists, for their own understanding of intelligence. If agility in adaptation to the environment is an indicator of intelligence in all living things (and in engineered inorganic things as well), then this criterion should be good enough for intelligent

problem-solving in AI too. This may be neatly transposed, in the schematics of artificial intelligence research, to the effort to quickly or cheaply reach the stipulated goal. Hence, feedback becomes a programmed repeated quizzing of the program's current state relative to the goal state in terms of the opportunity cost of taking any given path.

Marvin Minsky clarified the continuity of influence:

“ The era of Cybernetics was a premature anticipation of the richness of computer science. The Cybernetic period seems to me to have been a search for simple, powerful, general principles upon which to base a theory of intelligence. Among the ideas it explored were the following:

...Negative feedback: The psychological concept of goal was identified with the mechanism of setting up a generalized servomechanism to reduce the difference between an input goal parameter and an observed system parameter. This idea was exploited in various mathematical directions, but the secret of intelligence was not to be found in “optimal control” or similar knowledge-free theories. Nonetheless, the difference-reduction concept, reformulated in terms of a symbolic description of differences, finally became a key concept in AI in the General Problem Solver of Newell and Simon.” (285)

During the course of the development of GPS, a number of better search techniques for syntactic or ‘blind’ search, in which the factual specifics of the search space are not known, were developed. These included breadth-first and depth-first search, evaluation by means-ends analysis as mentioned, and various subsidiaries of the latter such as the A* algorithm. The first two mentioned consist respectively of a shallow examination one or two nodes deep of each of the possible solution paths, and a search of each of the nodes in depth, until it fails or succeeds (286). In either case, these methods may be quite computationally expensive. The sundry algorithmic syntactic methods employing variations on means-ends analysis- best-first, ordered search, and A* algorithm- differ from this, being more cognizant of computational cost and hesitant in that proximity to the goal is evaluated at each node and the node itself pruned (not searched at all) if it is not promising. (Of the latter derivatives of means-ends, some were developed by programmers other than

NS, but obviously influenced by their initiative) (287). These techniques are blind but not dumb: the only technique so dubbed appropriately is "brute force" search. In such a search, the inventory of all possible actions is traversed without any calculation. Because the test problems were so sparse on facts, the emphasis was on the procedures, which were more significant for this early work anyway.

The advances achieved in the form of GPS consist of most of the skeletal concepts in Classical problem-solving, and particularly in search. In this unintended context, the GPS ultimately did serve the purpose of General Problem Solving. It proffered at least indicia of means to solve problems in specific, as well as the apocryphal problems in general. The concept of the search space as metaphor for contents of memory and knowledge was elucidated and greatly improved under the aegis of solving specific problems in specific subject matters. The concept of simply deriving some systematic approach to search in any form was itself extraordinarily novel. The earliest development of the structural components of the field may be seen here- search, state spaces, goals and operators, the definition of problem spaces and the problems of blind search. The fact that these methods were later rendered outdated does not obviate the force of the achievement- in this specific case or in general.

Computational Chess-Playing

Chess, as we saw, is a tremendously stimulating topic for early AI. It continued to be important. We saw in the earlier chapter about the Rand Corporation that Newell had become interested in using rules of thumb as a means of developing computational chess. But by the mid-1950s, Simon was interested as well, and McCarthy was engaged in discussions with Arthur Samuel about his checkers-playing program. The NSS program, which was completed in 1958, indicated significant success in this domain. Simon had become interested in the field as early as 1952, when he wrote an appendix on the topic to one of his papers, in response to a summer seminar on chess-playing given by Von Neumann. When the paper, a famous economics work called "A behavioral model of rational choice", was published in Models of

Man (1957), the editors declined to include the appendix, apparently over Simon's objections (288).

Newell and Simon's chess papers, together and separately, were substantive (289). Their 1958 program emulates the other most notable ones of this decade by implementing static evaluation at the dead position, as suggested by Turing; a move generator based on substantive goals (the safety of the king, control of the center of the board, etc.) (290). It differs by using a number of specifically heuristic measures. It uses a move generator, which balances programs against goals, and compares alternatives. Bernstein's program had had a move generator as well, but used an additive evaluation function rather than heuristics. This is something that the other contemporary programs had not done. Even by 1958, the team had only programmed (or coded as they say), the first two goals (291), but both the technical or programming and the intellectual features of the game held their interest. Later in the 1950s, Simon became very much concerned with the psychology of chess-playing protocols, turning to De Groot's study of skilled players and then trying to reiterate them. Eventually, chess playing using computers became as much concerned with the sheer volume of computing power available as about the heuristics of top game strategies, but in its early phase the advantage was definitely accorded to heuristics.

The difficulty of the task did not deter Newell and Simon from lightheaded and optimistic statements. Their 1958 paper in Operations Research featured one prediction that undoubtedly seemed excessive at the time:

"...That within ten years a digital computer will be the world's chess champion, unless the rules bar it from competition". (292)

This giddy claim was used against them by AI's detractors. But the adage that he who laughs last, laughs best is illustrated by the fact that the prediction came true in 1996-1997, when the chess-playing program Deep Blue, product of joint research by CMU and IBM scientists, won a game in a series against Garry Kasparov, the reigning world chess champion.

The ultimate shape of the emulation of intelligence was undetermined. This is quite evident when we look at other research being done during the late 1950s. The theorem of

intelligence as primarily manifest as a neurological phenomenon of self-organization at a local level remained influential, as evidenced in the papers of Minsky and McCarthy as will be discussed in Chapter 11. Moreover, AI research often concerned pattern recognition rather than knowledge representation. This preference prevailed for nearly another decade. In the last chapter, we reviewed the early form of connectionist theories of intelligence. These persisted through the 1950s, but were drained of intellectual resources and finally dwindled away in the 1960s.

Chapter 11. McCarthy and Minsky begin Research at MIT

1. The Origins of AI at the Massachusetts Institute of Technology

The different intellectual agendas of different aspects of early AI were carried out by the establishment of distinct institutions. As we saw in the previous chapter, Newell and Simon proceeded with their opus, as a pair and with others, at the Carnegie Institute. Minsky and McCarthy both moved to MIT, the fervid epicenter of the Cold War's work. It is hard to imagine a better fit between intellect and institution. Minsky would stay for his entire career; McCarthy only until 1961, when he left for Stanford. During the first few years that both were assistant professors at MIT, they were phenomenally productive- creating the Lisp computer language, establishing the Artificial Intelligence Center, envisioning timesharing, and beginning to supervise dissertations on a wide variety of topics. The mutual capacities of Minsky and McCarthy, combined with MIT's receptiveness for brilliance and innovation, appears to have worked perfectly.

MIT, Home of Cold Warriors

Even today, brick and mortar universities are the rule rather than the exception. In the Postwar world, more than half a century ago, our robust and overwhelming online universe did not exist. Telex and telephone were the closest thing to virtual or online reality. Brick and mortar were the only possible venues for academic enterprises.

Even in the early years of the Cold War, the pursuit of hard science and weapons-related research was both MIT's past legacy

and its future aspiration. MIT, along with the Carnegie Institute of Technology and later Stanford, would be one of several original homes for AI for decades. It is not surprising, given its identity at the fulcrum of science and engineering. Cambridge was the original location for military-industrial science in the 1940s, and by the mid-1950s the actors, stage sets, and institutions for the next act were in place. Despite Harvard's advantage in computing during the 1930s, that institution was eclipsed as a center of computing hardware during the Second World War, and as the center for software development later on.

The Institute has been central to the history of computing for at least half a century; it has been an endlessly energetic and rigorous technological leader for longer. MIT was founded in 1862, and began to admit women as early as the 1870s. For the better part of a century it taught engineering as simply the reengineering and incremental improvement of existing machines (293).

During the Second World War, MIT's identity changed. The National Defense Research Council heaped funds upon MIT and particularly upon its famed Radiation Laboratory, or "Rad Lab". Stuart Leslie states that:

"...at the end of World War Two, MIT was the nation's largest non-industrial defense contractor, with 75 separate contracts worth \$117 million, far ahead of second-place Caltech (\$83 million) and third-place Harvard (\$31 million)". (294)

Following the end of the war, MIT redoubled its developmental ambition (295). The research "troops" were not demobilized, but were redeployed to various places on campus (296). These included the Lincoln Laboratories and to greatly strengthened MIT departments. Some laboratories were simply renamed: the Charles Stark Draper Laboratory was born as the Instrumentation Laboratory during the war, as was the Research Laboratory of Electronics, which was the renamed "Rad Lab". The previously unheard-of memory and storage requirements of Project SAGE brought previously unheard-of torrents of financing to computer hardware. The tremendous largesse also resulted in the foundation of "hard"- that is, highly quantitative- social science departments including those in linguistics and psychology, and in the bolstering of the Institute's pure science facilities (297).

Through the 1950s and beyond, the Institute became even more intensively intermingled with military contracting: by the beginning of the 1960s:

“ its contracts with DoD totaled \$47 million, plus additional obligations of some \$80 million to its federal contract research centers, Lincoln and Instrumentation Labs... prime military contracts for 1969 topped 100 million...” (298)

Reflecting finances for MIT proper, this figure does not even include MITRE, the Air Force contractee which had been spun off from MIT in 1958 (299). Leslie points out that MIT became so industrialized, and so devoted to constructing weapons, that its identity seemed itself confused:

“ Sizing up MIT in 1962 from his perspective as the Director of the Oak Ridge National Laboratory, physicist Alvin Weinberg, who coined the term ‘big science’, quipped that ‘it was becoming increasingly hard to tell whether the Massachusetts Institute of Technology is a university with many government research laboratories appended to it, or a cluster of government research labs with a very good educational institution attached to it.’ With nearly 100 million dollars in annual government-sponsored research contracts by the early 1960s (a figure that would almost double by the end of the decade), science and engineering at MIT had become big business.” (300)

As at Stanford, although not at CMU, computing started as a service department, rather than as an object of academic study itself. MIT’s Computation Center was founded in 1956, with money from the Office of Naval Research and an IBM 704 computer. It was intended as a service bureau for thirty universities and centers in the Northeast: scientists would present programmers with technical problems requiring punishingly long calculations, and the answers would be provided. The problems included:

“...fallout radioactivity in rainwater, a dynamic model of competition between two firms, a heuristic strategy for computer game playing, United Nations office operation, shop motions in irregular waves...”(301)

In its first few years, the Computation Center’s major concern was with the improvement of the computing infrastructure- which concern led, in turn, to McCarthy et al.’s timesharing (302).

Whatever its merits, in the context of its times, this did not allow for hacking or student usage of computers.

MIT's computing community, which would later become inextricably involved with AI, reached well beyond the university itself. Several laboratories with computing facilities have been mentioned. A description of the community is also incomplete without mention of BBN, a spinoff of several MIT engineers who built a consulting business. Bolt, Beranek and Newman was closely aligned with Harvard and MIT faculty members and graduate students in computing sciences for most of the half-century duration of the Cold War. MIT professor and IPTO director Robert Kahn called it "the cognac of the research business" (303). A number of the AI researchers mentioned in this chapter and the ones to follow were employed at BBN, which used AFOSR funds to subsidize their graduate work (304).

More than the other two major AI centers, MIT's institutional involvement with the Cold War contributed to its development of AI in many auspices. MIT had been a technology center, surrounded by other technology centers, for decades before AI got started. Due to its intimate involvement with the Cold War, the institution took on a Postwar structure very early. The design of the Cold War university or 'multi-versity' structure itself turned universities from liberal arts colleges into massive, multi-faceted institutions. These were home to huge numbers of permanent non-faculty sorts- i.e., research people, tinkerers, and gadget builders. More people who got started with their questions with Cybernetics, or with radar and the early Cold War technologies- radar, acoustics, electrical engineering- strictly speaking, would lend a practical and engineering tone to the AI done at MIT. This helps to explain the profusion of artificial receptors and effectors, and people who pursued this research into AI. Thus, there was a good deal of "bottom-up" AI - creating intelligence from embodiment toward mind. The purportedly bottom up approach- through tinkering, not with the economic man and decision making on IBM punch-cards, did indeed lead to the top-down questions, in robotics and in vision.

Marvin Minsky arrives at MIT

Artificial Intelligence at MIT developed at first from the bounty spilling from this set of historical circumstances. Marvin Minsky became a staff member in the Research Laboratory of Electronics, and taught an MIT course on automata and Artificial Intelligence (305). He finished his Junior Fellowship at Harvard, and then joined MIT's Lincoln Laboratory, working in the research group run by Oliver Selfridge. Despite the initiative of the theorem-proving algorithms that Minsky had distilled from Euclid, he did not work on AI per se for the next two years following the Dartmouth Conference in the Summer of 1956. He did, however, invent a device to measure human vocal pitch (306).

Minsky also began to write a survey of the field of AI at the time. This paper was repeatedly polished until it had become a long monograph, which was finally turned over to Edward Feigenbaum and Julian Feldman for their 1963 survey of the state of the art of AI. Minsky did not raise his own 'family' of graduate students until slightly later, when he was at MIT. While the first dissertations produced by the NSS students appeared late in the 1950s, the first round of MIT protégé dissertations began to appear after 1960 (307).

Minsky's brilliance inevitably produced both light and heat. Minsky himself did not follow up on the concept of theorem proving of geometrical figures after the conference, but which the proof of Euclid's theorems inspired Oliver Selfridge to pursue automated theorem proving. His 1956 notes on theorem proving were also taken seriously by Herbert Gelernter, a new physics Ph.D. hired by IBM as a researcher in the Theory of Automata Group in Poughkeepsie. Gelernter, along with Nathaniel Rochester, developed working computer programs which could prove geometric theorems, extending the initiatives of both the Logic Theorist and Minsky's ideas (e.g., Gelernter and Rochester 1958) (308). John McCarthy also contributed to this early work in theorem proving and heuristic programming, which by default was obliged to contribute significantly to programming languages for AI. The effort at IBM was concluded in 1959 when the company withdrew from several areas of AI (309), but the field itself continued.

The Invention of LISP

John McCarthy returned from the Dartmouth conference and the IRE Conference to his post as a Dartmouth math professor- but not for long. MIT recruited him that year, and a year later he moved to Cambridge to work at MIT's Computation Center, as a Sloan Foundation fellow, funded by Dartmouth College (310). McCarthy joined the Mathematics department at MIT in 1958, and he and Minsky established the Artificial Intelligence Project at the start of the 1958-1959 school year (311).

During these years, McCarthy spearheaded or single-handedly introduced three more important innovations in computer science and AI. Timesharing, a convention which permits both closer interaction with the computer program in progress and more people to work on a program at one time, will be discussed in later in this Chapter. Consideration of the 1959 Advice Taker program, a new architecture for AI and precursor to McCarthy's following decades of research in logic, follows this segment. The next several pages will concern the LISP computer language, a unique achievement on McCarthy's part.

As we saw in the discussion of the Dartmouth Conference, computer languages had been among several fundamental issues in McCarthy's mind throughout the decade. He refers to languages in the proposal that led to the Dartmouth Conference- ("During next year and during the Summer Research Project on Artificial Intelligence, I propose to study the relation of language to intelligence..." McCarthy et al. 1956). Moreover, in 1955 and 1956, he had talked about the topic at length with Allen Newell and Herbert Simon. During his time at Dartmouth, McCarthy had worked with mathematician John Kemeny. Kemeny and others later invented BASIC, the 'beginners'-all-purpose-symbolic-instruction-code' (312). BASIC was an important computer language for very small memory computers- which was all that existed at the time. During the summer of 1958, McCarthy conducted research using the mainframe computer at IBM's Poughkeepsie research center. There his ongoing pondering of computer languages resulted in the invention of the LISt Processing Language, or LISP (313).

LISP was not the first list-processing language- that honor belongs to the Information Processing Language (IPL). But Lisp did become a very widely-used vehicle for Artificial Intelligence

research for at least twenty years. In order to grasp its importance, we should consider what makes LISP different from other computer languages. This in turn leads to defining what computer languages are. A computer language is a means for the ordering and processing or manipulation of data. The best means by which to do this necessarily depends on the objective for which the data is being manipulated. Widely used business-oriented computer languages such as COBOL relied on a strategy entirely opposite that of LISP. The procedural languages as they are called, carry out particular procedures. They specify variables which then undergo a series of precise routines. The latter are not particularly changeable, while the former variables are strongly typed (that is, are not malleable with regard to the form of information being taken in). These procedural languages, which ordered the entry and processing of data in a certain fashion, using a specified order, are also called 'imperative' computing languages). Such programs provided exact steps regarding how to proceed for every programming task.

A tool is a good tool if it fits the usage to which it is being put. Under this criterion, procedural languages were appropriate for "mass production" data processing. The need for congruence of the language and the concepts being expressed is probably met with such tools. But they were not adequate for really addressing symbolic programming. Expressiveness, and extensibility, or amenability to the expression of new data structures and syntactic operations, are key goals for functional computing languages. Because the central idea of a functional computing language is that it be able to perform or apply certain functions, specified by the programmer, upon the data, these are also called non-procedural, or applicative, languages. This sort of computer language, in the form of Lisp and OPS (the successor for IPL at CMU) became the one most favored for AI.

A more specialized language was needed for AI, and LISP included from its origins a number of features that met this demand. AI demands that ideas be defined in more than one context, that things be defined elaborately, and that opportunity be provided for the repeated introduction of new categories to existing classifications. Several "revolutionary" features of Lisp meet these needs (314). First, Lisp supports highly nested

structures. That is, in defining an object, Lisp allows for both the definition and the practically infinite re-definition of that object, using a series of parentheses. Second, Lisp is "extensible", or amenable to alterations and to the creation of new programming tools or types of data. One may specify the formula for an abstract concept in advance, without having the data per se. Moreover, LISP is highly "function-oriented", allowing for creation by the programmer of new routines for addressing data. This feature sets a striking contrast to the organization of computer languages as procedurally rigid hidebound routines prevailing in the late 1950s as just discussed. Lisp also has a specific variable type called the 'lambda function' which is used for the definition of novel variables (315). The ease of conditional statements representation makes the construction of delicate hierarchical distinctions between different kinds of objects simpler.

A description from the bottom-up is more appropriate than one from the top-down in this case. The language is elegant in design in both the aspect of its rudimentary particles and the potential complexity which can be built up from such elements. The "particulate" or rudimentary element in LISP is the atom. The atom, or atomic formula, is simply a data type which is infinitely definable and redefine able. Anything may be placed in parentheses and defined as a data entity, put into a function, and hence placed into the programmer's opus. Once labeled, atoms may have properties associated with them: these are referred to as property lists. "Lists", a slightly larger unit, are similarly malleable. In the case of each feature mentioned above, the emphasis on abstract features meant that the language's activity could be defined by the programmer rather than dictated by a preordained protocol (316).

This freedom of choice in program design is the basis of non-procedural languages, a major innovation of the 1960s. McCarthy was also closely involved in the design of ALGOL, an influential invention for numerical computing, during the design of LISP in the late 1950s (317). Endlessly tested and refined by hackers at MIT's AI Lab and later at Stanford SAIL and most other AI institutions, LISP greatly facilitated the progress of AI.

Theorem Proving and the Advice Taker

John McCarthy had invented the name 'AI', but was not the exclusive inventor of the substance of AI as cognitive emulation. But during the remainder of the 1950s he was quite a homesteader, so to speak, in terms of the development of the field. He worked intensely on the building of computational infrastructure, in the form of both the LISP computer language and in timesharing. He also began to work on formal logic late in the 1950s.

When they started exploring computational problem solving, NSS had made their extremely difficult work somewhat easier by limiting the terrain of both problem or goal and area in which the goal can be found. They started with a relatively small circle, as in the case of the Logic Theorist. Alternatively, they also used a small number of canonical devices for limiting search (as in the case of the chess-playing programs). They also made a clear discretionary decision to model input and output in ways that did not resemble how people functioned: that is, both IP and OP were in computer languages, on paper punch cards. John McCarthy, on the other hand, began to focus on the development of formal languages which would be able to push any problem back into a tiny circle. By commencing work on the formal system of expression rather than on the agent's navigation of the environment, McCarthy took on a task that the logicians had not yet arrived at. Moreover, his task could not take cues of the same sort from cognitive psychology, as NSS could, and by not limiting the semantics of the things that the program could be told, he was coming up against some very knotty problems in AI.

The heroic, but initially impractical step of trying to make the program draw the circle around the world, rather than limiting the data given to the program, began for McCarthy with the Advice Taker. This was a proposal for a program, which McCarthy and Minsky were planning to write, which would solve problems that seemed more in line with everyday human intelligence than the NSS problems. In 1958, McCarthy presented the paper 'Programs with Common Sense' at the Symposium on the Mechanisation of Thought Processes in England (318). He professed dissatisfaction with the existing work which solved problems in logic or algebra, which were contrived or artificial (RES' terms) relative to the sort of things that people do on an everyday basis. He offered instead

criteria for “intelligence of human order”. He proposed a wider, and in a sense more ambitious criterion for intelligent programs: he wants to have a program that solves ‘common sense problems’ by using deductive logic, that is figuring out things like logistics based on a set of formal statements about the real world (e.g., one often uses a car to drive to places). The common-sense problem that he offers is meeting a schedule which requires arriving at the airport on time (319). In real life the solution to such a problem might, under the best of conditions, be relatively straightforward to anyone who would be required to show up at the airport. But that is only because we have picked out the solution to the challenge from a multitude of other items which might or might not be relevant to getting there- the weather, the road conditions and traffic, the person doing the driving, the timing of all parties involved, the availability of the car, other means of transportation. When one starts to pull the problem apart in this way, the problem of getting to the airport stops being so simple that it can be solved by a program which takes imperative sentences and ‘knows’ how to represent the semantics of the sentences to solve given common-sense problems (320).

The paper was criticized ferociously by Yehoshua Bar-Hillel, who in the discussion after McCarthy read the paper called it “half-baked” for assuming that the process of deduction was an adequate framework for the behemoth of human common-sense intelligence. Certainly, the original presentation is sketchy- but for the right reasons, that is, because McCarthy was beginning to grapple with big questions. McCarthy introduced several significant problems, which Bar-Hillel properly identified and which have occupied McCarthy and others for decades thereafter. In a sense, common sense would prove to be the elusive Ultima Thula, the legendary land of Northern European mythology which remains impossibly far-off and difficult to reach or even envision. The Advice Taker offers to design a program with common sense. This grand initiative was substantiated by AI work in naive physics, that is, the real knowledge of the world that common sense embodies. The proposal also provoked the first statement of the intractable issue known as the Frame Problem, that is the issue of framing, and therefore getting a grip on, the material at hand amongst the multitude of things that are in one’s field of

vision. The first person to actually say it seems to have been Bar-Hillel, who pointed out that McCarthy's proposal not solved the problem of picking out the relevant needles from the haystack of facts in the world:

"... I do not think there could possibly exist a programme which would, given any problem, divide all facts in the universe into those which are and those which are not relevant for that problem. Developing such a programme seems to me to be by 10 to the 10 orders of magnitude more difficult than say the Newell Simon problem of developing a heuristic for deduction in the propositional calculus. This cavalier way of jumping over orders of magnitude only tends to becloud the issue and throw doubt on ways of thinking for which I have a great deal of respect. By developing a powerful programming language you may have paved the way for the first step in solving problems of the kind treated in your example, but the claim of being well on the way towards their solution is a gross exaggeration. This was the major point of my objections." (321).

The frame problem, as Bar-Hillel points out, is a significant issue which is not solved simply by better programming languages. Any extension of formal logic, for instance writing first order predicate calculus as a programming language, must turn philosophy into knowledge engineering. Finally, the Advice Taker paper inspired McCarthy's work on circumscription and formal languages. The Lisp language was intended to provide a superior way to express statements about objects in the world of an AI program: McCarthy and his colleagues have spent decades since then in improving formal logics and associated computer languages so that they can express time, sequences of events, modal or subjunctive logics (that is, 'the road not taken'), etc. While the Advice Taker itself was never written, McCarthy has spent decades since 'baking' the ideas contained in the proposal.

The Birth of Computer Hacking, and Its Rewards

hack: "... 2. n. An incredibly good, and perhaps very time-consuming piece of work that produces exactly what is needed.".. 6. vi. To interact w a computer in a playful and exploratory rather than goal-directed way... 9 [MIT] v. To explore the basements, roof ledges, and steam tunnels of a large institutional building to

the dismay of Physical Plant workers and (since this is usually performed at an educational institution) the Campus Police...” (322)

Practically all cultures- nations, ethnic groups, artistic genres such as theater and painting- have both a ‘high’ tradition and a ‘low’ one. For instance, Latin is contrasted strongly with the vernacular languages of Early Modern Europe. Strindberg and Ibsen are dramatic statements that have moved many people, but so are ‘Seinfeld’ and ‘The X-Files’. The low culture invariably gets a bigger audience; it wins the popular vote.

And so it is with artificial intelligence. All of the major figures, who founded laboratories and research programs have played the role of the leader of high culture. But all have also presided over the folk tradition of AI in their laboratories. The functional equivalent of low culture in AI is hacking. Like low culture itself, it is viscerally appealing, useful, interesting and not at all profound, whatever that term means. It is often small in scale (at least initially), or calls upon few resources to create it, and is often made by people with relatively less formal education than the adherents of high culture.

Hacking, the folk culture of AI and computing more generally, consisted of work done largely by undergraduates and others not particularly cognizant of their place in the graduate school-to-professor pecking order, and largely outside the auspices of academic progress. Some people at MIT, and their work, have moved along the regular academic track- Gerald Sussman, initially a vision researcher, for instance, has become a renowned professor. But many hackers seem to have been dropouts, and the work contrasted with the formal culture in the context of which people worked for grants and other things of the academic course. At the same time, MIT is proud of its tradition (323). Likewise, much of the work done in hacking has concerned the emulation of intelligence- in vision, robotics, ergonomic utilities which facilitated the cumbersome features of computing and applications. But this work has generally been not much concerned with theories of the mind as so much of AI has been, and has not had an academic axe to grind (in either a positive or a negative sense of this phrase). Hacking is sometimes defined as learning about computing by trial and error, or usage of

computers in a random way or for fun, or for enjoyable projects rather than some scientific purpose. The same sort of tinkering can take place upon other objects of study, in areas other than computing or AI.

One other issue is important to note: hacking was not, at this time, in any way associated with financial fraud, theft, tampering with military or corporate computer systems. As Oliver Selfridge points out:

“ Remember that the term "hacker" back then included a tone of admiration--someone who could sit up all night getting something to work. It had none of the corrupt implications that it has today. Hackers in fact, as I remember, were largely responsible for the amazing advances which triggered everything in computers.” (324).

Computer hacking thus has a pre-history, just as formal AI as the physical symbol system or different forms of representation of memory has a prehistory.

We have established that AI as high culture was barely extant in the 1950s, with the exception of a very few people such as Newell, Shaw and Simon at Rand and CMU. But there was plenty of pre-AI, in the creation of the Perceptron, among the Cybernetics group, in the early work of the cognitive scientists, in the theories of automata and self-organizing systems by McCarthy, Minsky, Selfridge and Von Neumann. This is the prehistory of formal AI. There is also a pre-history to the other sort of informal AI, better known as hacking. There is a significant hobbyist tradition in all electrical-related areas- radio sets, electrical wiring, model railroads, car mechanics, self-guided mobile planes, and the like. Much of this continued in force through the latter part of the Twentieth century, but was gradually eroded by the increasing popularity of computing, and its ready availability after 1985 or so. Popular Mechanics, a serial devoted to electrical and mechanical projects for home-garage hobbyists, enjoyed immense popularity in the United States during the PostWar decades. Electro-mechanical engineering hobbies, epitomized by model train clubs- MIT's being the most famous (325)- and other Popular Mechanics-types of tinkering, absorbed people who would otherwise, or later, become entirely entranced with computers.

This Ur-hacking was first turned away from electrical hobbies and toward computing in the tinkering of a dozen or so kids at MIT's Model Railroad Club (TMRC). In the absence of computer facilities accessible to undergraduate students, this opportunity evolved into elaborate programmable switching programs. Students were given complete control and as much time as they wanted at the Club's elaborate train system (326). Potential students lacked any opportunity to work directly with computers, and thus stuck to pursuits which could absorb as much technical ingenuity as they wished to throw at them. The first computer hackers were apparently recruited to, or drawn to by word of mouth, from this cohort, at the very end of the 1950s. Minsky built alliances with the students in the popular Model Railroad Club, where "indigenous computing" was already going on (327). John McCarthy, likewise, taught one of the first college courses in computing, and arranged to allow undergraduates computer time to punch and run their own code for the course. At CMU, students obtained usage of machines through graduate professors. Newell and Simon and colleagues achieved early results, it seems, because they arranged early computer access for their students. Ed Feigenbaum, for instance, took a course in IBM 701 programming with Simon in 1956 (328). Starting in 1959, McCarthy taught a course which required programming on the IBM 704, and offered CPU time for students. Minsky and his EE colleague Jack Dennis began to cultivate the friendship of a number of the students who had moved from MIT's model railroad club to the computer.

This was quite a revolutionary concept. The status of computing, at MIT and elsewhere, circa the end of the 1950s, was certainly better than it had ever been technically. But computers remained infinitely inaccessible to the average interested party. Until Digital Equipment Corporation was founded in 1957, IBM's rivals were considered paltry. The industry was characterized as "IBM and the Seven Dwarves", and thus the practitioners of AI and everything else had to accept the computing that they were given. The machinery for computing was, in academic settings, typically an IBM 704. The machine itself placed stringent limits on the nature of access to computing time. The IBM 704 required two larger rooms and constant monitoring in case the special air

conditioning broke (which happened often). Only people with some official usage, typically with some connection to Lincoln Laboratory or the RLE, were allowed to hand over punch cards to the machines systems operator. This effectively ruled out opportunities for curious undergraduates, at MIT and elsewhere (329).

Regular people, even MIT students, were no more allowed to work the controls of a mainframe computer than they can operate the controls of a nuclear power plant today. “Friendliness” to the end user may seem like a commonsensical notion to the reader at the turn of the twenty-first century, but in the middle of the twentieth, it was by no means obvious. Because of the sensitivity of the work done by these machines, and because of their very proneness to error and breakdown, highly restricted access and layers of guards were the rule here. We should also consider that the management philosophies of the time were still not terribly far removed from the Pre-WWI philosophy of scientific management, which purported to shield valuable machinery from worker ‘ineptitude’ or sabotage.

There was also very little pressure from users to help initiate further innovations in computing. The tiny number of people in the field did not exert a great deal of pressure on the demand side of the market. The miserly supply of computer time even for professors was commensurate to the generally limited demand for the machines. The stringent limitations of any individual’s time with the computer was due in part to the larger lack of active demand for computer time in all but the biggest corporations and most elite university-military settings. These two market forces, informally speaking, were rather sluggishly matched to each other at the time, and only altered a bit later in the 1960s- which we shall get to in a moment. As the economist might say, the market exhibited equilibrium at a low level. What need there was for change was not because of students but because of need for more programming time for programmers in scientific and business settings (sorry, hackers).

McCarthy and Minsky thus brought in their earliest students, including both scholars who completed Ph.D.’s and ‘programming bums’, now called ‘hackers’, who did not care enough about academic degrees to finish them. This neglect of credentials

represented a loss for the university Bursar, but not for technological progress. Hackers, with time to burn and enough brains for ten or twenty, have been perhaps the most prolific inventors of computing's process innovations (innovations which make existing things work better). The occasional exhibition of "hacks", meaning elaborate pranks for which MIT students are famous, has apparently been part of MIT's culture for decades. But around 1959, for some undergraduates at least, the practice turned to computing became more regular, as the work was first used for clever exhibitions of intelligence on a computer (330). For some it turned into a way of life, which superseded classes and graduation. Regardless of effects on academic transcripts, the hacks proved immensely productive for computing per se. MIT was preeminent in such adaptations, as much because of its unique local intellectual culture as because of its early contributions to AI.

It is hard to say which was more influential- the local culture of 'hacking', or tinkering with machinery, or the availability of the newest machinery from DEC, or ARPA's latitude in ascertaining intellectual freedom for hackers and their professors. In any case, this particular mixture proved uniquely productive. Note that hacking was really part of the productive AI environment. Much- perhaps too much- has been made of the culture and personal quirks of the hackers. We won't tarry on the topic; see Stewart Brand, Mark Levy, and Hapgood for excellent accounts. The philosophical base behind hacking was loose and intuitive rather than grand. A fine illustration, perhaps, is that of a famous 'hack', or prank, several years prior to hacking computing, in which a group of MIT undergraduates contrived to lift the complete body of a police car, with a functioning siren, onto the grand dome of the MIT campus (331). It was apparently near impossible, bound to impress, and above all, fun.

Thus much hacking was basic and everyday as AI's means and ends are complex. Unlike the police car escapade, computer hacking has bequeathed the world useful things. The first principle was that computer facilities should be readily available to one and all, and relatively easy to use. The philosophical axioms which justified hacks included a belief that "information wants to be free", as Ted Nelson said in Computer Lib. In certain

cases, it is difficult to believe that the instinct to make computing activity more democratic and freely available was not an echo of the Marxian axiom of putting the means of production into the hands of the forces of production- that is, the workers. This idea is bolstered by the fact that the movement to popularize computing apparently included a number of 'red diaper babies', that is, the children of Old Leftists (332). Thus, both the Leftist challenge to capitalist authority, or any single authority, over machinery and wealth, and the more simple American and Anglo-American tradition of tinkering with machines were embodied in hacking.

Hacking on train sets, radios, and remote-controlled boats and cars and the like had been the rage for much of the century. Hacking on computers began almost as soon as it possibly could begin. The aperture to slightly more wide access to computing began to open by 1960, and proved quite revolutionary. User-centered innovations such as games, on-line editing, debuggers, and even word processors, seem to have been in large part the result of computing time made available to university students. McCarthy and Minsky had offered students access to the computing terminals, in the late 1950s before timesharing really existed, and in so doing brought in a few adherents.

The Initial Invention of Timesharing

Timesharing would be central to all multi-user computing after 1960. But it began almost as soon as computers themselves, and it was named- again by John McCarthy- and greatly improved in the late 1950s. Timesharing, as we saw, was practically a necessity for AI- how else to allow endless hours of experimentation and code-writing. Timesharing works by intermittently doing bits [sic] of each allotted task until every task is accomplished. Technically speaking:

“ in a multiprogramming computer, several programs may be processed concurrently by switching from one to another in a fixed sequence to permit a certain number of instructions to be performed on each occasion” (Penguin Dictionary of Computers).

Or timesharing can be defined as:

“A system in which a particular device is used for two of more concurrent operations. Thus the device operates momentarily to fulfill one purpose then another, returns to the first, and so on in

succession until operations are completed”, Webster’s NewWorld Dictionary).

John McCarthy and a number of other people came to this conclusion early on in the game, and created timesharing, a programming artifact in which two or several, or well down the road, hundreds of users could program at the same time, each apparently given individual access to the central processing unit. JMC was not the sole inventor of TS, as we have seen, although he was apparently the first one who proffered the idea of a ‘computer utility’, or “a community utility capable of supplying computer power to each customer where when and in the amount needed. Such a utility would be in some way analogous to an electrical distribution system” (333).

The essential idea of timesharing, from the perspective of the processing unit, is the system interrupt, which stops one program in order to implement a second (or third, u.s.w.) computer job. The first implementation of this idea had been in the Atlas computer, which lacked commercial success or wide distribution. The Atlas was designed in Britain in the mid-1950s in cooperation between the University of Manchester and Ferranti Ltd. (334). In the Atlas, this programming design feature was called an “extracode” instruction, and it was intended to augment the internal memory, rather than to make programmers happy. The Atlas was intended to provide a large memory, with a space of 1 million words of 48 bits each. This was far too ambitious for magnetic core memories at the time, so the designers gave the machine a relatively small core memory (16,000 words) and a far larger rapid magnetic drum memory (96,000). A datum missing from the main memory’s “page registers” was sought by an interrupt, which stopped program execution while the page was moved from the high speed drum. The system interrupt was essential to the proper functioning of the virtual memory (335).

Thus the Atlas was intrinsically timeshared. The Atlas’ engineers originally included timeshared terminals in their designs. This idea was taken out of the blueprints as being too expensive- unfortunately, as Michael Williams’ A History of Computing Technology observes:

“ Had this been incorporated into the machine, we would likely have seen the mass produced time-shared computer being

commercially available a few years earlier than it actually was.” (336)

The system interrupt had not been introduced as a way to alter the nature of human-computer interaction, but properly implemented, this feature did just that. The essential idea of timesharing from the perspective of the user is many people preparing and editing their programs simultaneously, at consoles which combine teletype and CRT screen units.

But the demand for timesharing did not grow only from the side of corporate computing. Timesharing was a populist movement, borne by self-organization among programmers, as well. Late in the 1950s, the technical model of timesharing was independently conceived again by others. British engineer Christopher Strachey thought of the concept of timeshared programming through the addition of several individually operating and card-reading consoles adjacent to the mainframe computer (337).

John McCarthy's Innovation

In addition to Strachey and the Atlas designers, John McCarthy also apparently invented timesharing, including most of the modern I/O and programming features, on his own (338). He first tried to change the available computer- which was necessarily an IBM 704- in the Fall of 1957, when he arrived at the M.I.T. Computation Center on a Sloan Foundation fellowship from Dartmouth. He proposed a programming fix to said computer. The programming fix would allow an expansion of the buffer zone so that it could work iteratively between one program and another (339). IBM was persuaded to allow MIT to lease the machine, an early keyboard input device known as a Flexowriter, and work began to modify the machine. Oddly from the perspective of 2010, it took “a year, perhaps two” to actually receive the IBM 704 delivery. The idea was sufficiently technically possible and so helpful for programmers that both the supply and demand sides helped to make it happen. Originally, pleas to IBM resulted in hardware modifications to the IBM 7090, which MIT received around 1960 (340). Then the newly formed Digital Equipment Corporation was persuaded to modify their computers to allow a higher level of interactivity. DEC hired several MIT electrical

engineering graduates who were familiar with the equipment and the need for greater flexibility, and collaboration resulted in the new PDP-1. McCarthy's own programming employee Steve Russell began to work on this, and quickly was joined by several people who proceeded to make time-sharing a substantive part of their life's work.

McCarthy's proposal was taken seriously: MIT called a committee to study the topic. The Long Range Computer Study Group consisted of electrical engineers (including Jack Dennis, McCarthy himself, and Marvin Minsky). The panel concluded that time-sharing was a good thing and that there ought to be more of it on the campus (341). The work on timesharing may not have been taken as seriously as McCarthy would have liked: this is considered one of the reasons why McCarthy accepted Stanford's offer in 1960, and did not wait to hear what the Long Range committee had to say. MIT implemented a timeshared 7090 in the RLE, to the immense benefit of the early electrical engineering students and hackers who needed computing facilities. Next, the MIT computing center, which was the service center that took in computer jobs for other institutions as well as MIT, set up CTSS, another such system (342).

This was only the beginning, and the time frame stretches forward out of the purview of our study. The National Science Foundation would donate money when the work proceeded further. But it was ARPA's money, a few years later, that enabled MIT to put its money where its mouth was. Project MAC, funded by ARPA, allowed further computer purchases. Throughout the 1960s, timesharing would be implemented and then improved by successive MIT and associated projects- CTSS, the BBN project, MULTICS and finally Unix- over a period of decades (see Ceruzzi 2003, p155+). Timesharing was the technical foundation of all multi-user computer systems, and as such its importance for AI cannot be overstated.

Institutional Advances in Computing

By the end of the 1950s, circumstances regarding computing were changing at MIT. McCarthy requested further computing facilities, and promising new ways to facilitate computer access. The national insistence on more scientific research began to

recast computing as a necessary scientific activity rather than a luxury. This would reverberate to computing, both the existing sort of scientific problems and the efforts in List-processing languages and other forms of AI.

The Defense Department- the biggest 800-pound gorilla imaginable- wanted 'bigger better faster and more' computing. By all possible accounts, this was demand from above. However, the changes took place from more indigenous sources as well. Levy observes the gradual popularization of computing at MIT among the hackers:

" In 1959, a new course was offered, was the first course in programming that freshmen could take; taught by John McCarthy; and McCarthy's advocacy of AI was thought at the time to be very silly; CS did not officially exist, so McCarthy taught in the EE department; McCarthy had started a program on the IBM 704 to teach it to play chess. (343)

The participation of the eager undergraduates created demand among them for increased time and access- in the midst of severe limitation to both. The timing between card IP and OP was overnight at best; instructions had to be perfect. Finally, there was no possibility of actually getting near the machine (344).

As undergraduates began to realize that computing was the greatest potential hack ever invented, generals and deans wanted to throw money at this phenomenon. In 1960, MIT dean, and later president, Jerome Wiesner encountered McCarthy and Minsky in a hallway. He asked them what sort of facilities they needed. When they requested, modestly, only an office and keypunch and two programmers, he lent them the labor of six 'redundant' graduate students from the Research Laboratory of Electronics. The RLE students were supported by JSEP block grant, and thus could be put at McCarthy and Minsky's disposal (345). IBM supplied the equipment. The Office of Naval Research, under Martin Denicoff, also acted as a sort of ARPA before ARPA, providing other funds which bolstered the volume of research that had been done in robotics prior to the existence of the IPTO (346).

Larry Roberts, later Director of ARPA-IPTO, remembered computing time as being available only to a few people at MIT:

" Around 1960, [computer science] virtually did not exist as a subject, or as an activity. They did not have computers. The 704

was the Computation Center's only utility. If you were one of the circle of people that worked on WHIRLWIND, fine, you worked on WHIRLWIND. If you used the 704 you were probably in a white jacket and worked for IBM. Otherwise you submitted programs from somewhere and where no students had access to them.” (347)

Relatively shortly afterward, funding for computer facilities, with improved IP-OP technology, would appear in a vast flood. Project MAC, the mother of all time-sharing systems, would be initiated in 1962, and has basically continued since then. The Sixties at MIT would be an astonishing proliferation of computing creativity- a topic from which we must reluctantly turn away.

Chapter 12. AI Hype, and AI's Detractors Before its Time

As we saw in Chapter Seven, anti-machine sentiments are a constant historical theme. They persist as with the same intransigence as the desire to build machinery which embodies intelligence. AI galvanized and attracted such sentiments, which had a set and generally unchanging series of objections. The equal and opposite reaction to the very existence of computers would continue to speak vociferously in the late 1950s.

AI's Philosophical Underpinnings, Revisited

AI embodies the most optimistic claims of Western philosophy, belief in progress and in the possibility of technology. The gut truth of Western philosophy, as both its friends and foes agree, is the coherence of the world and of the mind that perceives it. Empiricists Hume and Locke may have disagreed with Rationalist Descartes on many things, but the sublimity of the mind was not one of them. As an heir to the both philosophical traditions, AI subscribed to the Cartesian fascination with the mind and the Empiricist trust of the evidence of the senses. Certainly during the first years of AI, critics may have correctly discerned an unvarnished trust in the potential of science, AI included, to uncover the truth and build great things.

All well and good, but hubris is dangerous if not tempered by gravity and humility. The high and mighty must publicly profess

humility today, but none of these post-60s artifacts existed in the 1950s. The trait of humility was certainly not markedly evident among the scientific and political elite in the Postwar decade. Certainly the 1950s was one of the highest points of Classicism in Western civilization, and this theory merged seamlessly into the belief in American superiority among the scientific and technical elite.

If AI's very aspirations carried with them some of the hubris which abounded at the time, this is no surprise. But even amidst the delirium of science at the time, AI set a high water mark. By setting its metaphysical grasp and its technological sights higher than anyone else, AI inevitably established itself as a lightning rod to which friction would flow. It practically invited such attention. Moreover, insofar as the near-belief in omniscience in Western science was itself attacked, so too would the audacity of AI be attacked. AI's practitioners have been hardworking, nose-to-the-grindstone scientists, doing the sort of work that keeps people realistic and reasonable. But the support of the field by the powers that be meant that those critical of 'power sciences' criticized AI as well.

The logistics of the situation suggests predictable sets of friends, and enemies. Friends included the cognitive psychologists who worked with Newell and Simon almost from the moment they got started, and who both men cultivated at Rand Corporation's summer sessions and wrote with at the Carnegie Institute of Technology (348).

One would expect engineers and scientists to manifest enthusiasm for AI, but this was not uniformly the case. A number of Cyberneticists, such as Rand engineer Richard Bellman, attacked AI as fraught with hubris (349). J.C. Shaw collected anti-thinking machine quotations, some of them from Cyberneticists whose nerves AI had pinched. These statements ranged to the very extreme, for instance, as the British Stafford Beer who referred to computers as "less than morons".

Even to people inside the scientific community, AI seemed impossible, and somewhat illicit. The Rand Corporation's 1963 monograph celebrating its first decade and a half discussed Rand's analysis of nuclear warfare, its studies of "the economic and military capabilities of the Soviet Union", desalinization of sea

water, and linear programming. However, it did not mention either AI or complex information processing among Rand's achievements (350).

This book has perhaps emphasized the support that the fledgling AI community received, but the field's experience within the larger cohort of technical and scientific fields was not uniformly one big support group. According to Minsky:

“ The thing about AI was that nobody believed it could be done. Our worst enemies were other computer scientists. A lot of people thought that it would be humanists and artists who were skeptical, but actually it was always people down the hall.” (351)

Speaking in a 1989 interview, Minsky was referring to computational conservatives- so to speak- at MIT. This is worth noting this- Minsky's work took place at the United States' most preeminent engineering school. As we mentioned in Chapter One, and in Chapter Six, suspicion of higher claims for computers was rife at IBM as well, and on the part of some academics throughout the 1950s and into the Sixties.

One would anticipate that certain academic cohorts would be far more skeptical, simply by nature of their professional enterprises. A prospective list of AI skeptics might include humanists and artists, and especially philosophers. Was some of this sheer professional jealousy ? AI addresses many questions of philosophy, such as the nature of learning and perception, in a way that is bound to attract more attention than philosophy does. This is the same mundane reason that the humanist often considers the engineer unappreciative of the arts. The objections are more profound and personal as well. Engineers do not put humanists out of business, since engineering technology does not make the study of the humanities irrelevant. But the relation between AI and philosophy differs here. If AI fulfills the hopes of philosophy by turning its speculations into engineering, it also supersedes the claims of philosophy by literally turning philosophy into an empirical science. AI is an engineering field, which entails doing something rather than simply arguing about things one can't prove. Once that is done, what will philosophers have left to do ? As we will see, this quandary was not lost on certain philosophers.

Objections in the late 1950s

As we saw in Chapter Seven, objections to computers apparently appeared several minutes after computers themselves. Alan Turing described and safely defused a number of such objections in 1950. This certainly did not stop objections from appearing. An extreme example of such stipulations is the work of Mortimer Taube, which lumped AI together with some of history's greatest scientific impossibilities:

" 1) is it possible to translate by machine from one language to another ?
2) is it possible to build a perpetual motion machine...

...

4 is it possible to see God;

5 - is it possible to have extrasensory perception ?

...

10 - is it possible for a machine to think ?." (352)

These arguments are, for the most part, easily answered. This would not be true of the less superficial and thoughtless arguments forthcoming in the next decade.

A strong note of dissension concerning continued research in computing came from Norbert Wiener. Wiener's argument, in turn, was an echo of the past, for what he said repeated the assertions of Lady Ada Lovelace, one hundred and something years earlier. 'Lady Lovelace's objection' (1842), found in her notes to the de Menabrea lecture transcription, asserts that, "The Analytical Engine has no pretensions to originate anything. It can do whatever we know how to order it to perform (her italics)." (353) Turing had pointed out that this is, practically speaking, not true for a rather mundane reason: "Machines take me by surprise with great frequency. This is largely because I do not do sufficient calculation to decide what to expect them to do..." (354)

Wiener used a new version of the slavery-dumb machines argument, asserting that in the nuclear age this was something more frightening than reassuring. He proffers that intelligent computing programs are a moral problem as well as a security risk:

" The problem, and it is a moral problem with which we are here faced is very close to one of the great problems of slavery. Let us grant that slavery is bad because it is cruel. It is however self-contradictory and for a reason which is quite different. We wish a slave to be intelligent, to be able to assist us in the

carrying out of our tasks. However, we also wish him to be subservient. Complete subservience and complete intelligence do not go together. ...if the machines become more and more efficient and operate at a higher and higher psychological level, the catastrophe foreseen by Butler of the dominance of the machines comes nearer and nearer.” (355)

One of Norbert Wiener’s best qualities was his brave willingness to criticize his own work (e.g., in military contexts). The speculation as to machines overtaking human masters can’t be discredited because its time horizon is not specified. But practically speaking, he was not addressing the near future.

The Appearance of AI Hype

If AI took on the highest of aspirations, it did so with a remarkably light step. The journal articles of the early researchers were not spuriously marred with grandiose statements. As we mentioned earlier, Herbert Simon announced to his students at GSIA in January, 1956, that “Over the Christmas holiday, AI Newell and I invented a thinking machine” (356). Immoderate claims were made only very rarely. Despite critics’ assertions to the contrary, it was a rare occasion on which AI’s computer scientists said anything to justify the Dr. Strangelove nightmares. Twice in the late 1950s, AI practitioners made precipitous statements which were at best provocative and at worst ill-advised.

Slips by Newell and Simon

In 1958, Allen Newell made a speech (coauthored with Simon), to the Operations Research Society, predicting AI chess-playing and a high level of artistic creation within the next ten years (357). This prediction was quite prescient, as these things did indeed come to pass- although not in a decade. Yet when Newell and Simon predicted in 1958 the appearance of championship playing chess and original AI music within a decade, they elicited derision.

Rand researcher Richard Bellman responded angrily in the next issue of the journal Operations Research (358). He very nearly accused Newell and Simon of being fortune tellers who adhered to the “principle of optimism”, “Never make negative predictions.” The folly of predictions which never even come close

to realization is cause for laughter, he points out. But those who make predictions are practically indemnified, since "One valid prediction completely overshadows the one hundred wild ones." This accusation is a slur on Newell and Simon's honor as well as on their identity as scientists. Computer chess did indeed succeed, but more slowly than anyone had thought it would. Simon later stated that he had thought that many more people would enter the computer chess field than did during the early 1960s (359).

Simon was typically given to much more moderate statements in the academic scholarly press. For instance, in "A Computer for Everyman" in The American Scholar (360), he emphasized the continuity of human life rather than the revolutionary quality of changes wrought by the computer. These sorts of scholarly articles are of course ignored by news reporters and professional philosopher-critics.

Minsky Coins a Controversial Phrase

Minsky's affiliations and obvious brilliance garnered him ongoing invitations to conferences. In 1958, the first quotations by Minsky were sought by a press curious about computers, or 'mathematical machines' and 'electric brains'. This is the term that Newsweek used to refer to computers in throughout the 1950s. At the 1958 Mechanization of Thought Processes Conference, covered by Time magazine, he stated this openly:

" Many of the speakers tackled the question "What is intelligence ?" None of them had a wholly satisfactory answer. Dr Marvin L. Minsky of MIT felt that "the problem is unduly complicated by irrational human reverence for human intelligence. "We can often find simple machines, he said, which exhibit performances that would be called intelligent if done by a man. We are understandably, very reluctant to confer this dignity on an evidently simple machine."

Dr. Minsky is convinced that there is nothing special about intelligence or creativity. He thinks that as machines are built to perform more complicated mental processes they will gradually acquire more of the creative abilities of the human brain. When the first intelligent machines are constructed, suggested Minsky

(perhaps joking only slightly), they may refuse to admit that they are machines at all.” (361)

‘Meat machines’ used to refer to human intelligence, is actually an anomalous phrase given the fascination and high regard which AI practitioners hold for thought. While it may have been intended to provoke, it instigated outrage as well. It would be roundly criticized by Minsky’s colleague Joseph Weizenbaum about a decade later.

Digressions and Failures

If the road to Hell is paved with good intentions, then conversely- the road to success is paved with failures. The apparent success of the larger venture of AI was accompanied by a number of protracted failures. The historical oddity of the idea of AI did not lend it credibility early on. Instead it seemed to many people like a crank enterprise, and even rather sinister. The personalities involved in AI, most of them as wholesome and cheerful as the Good Humor man, have never justified any such suspicions. But the field drew both the latter suspicions and hopes for miracles. Moreover, some pursuits that were fraternal to AI, and especially early machine translation- crashed spectacularly.

Dead ends and false starts are the stuff that journalists covering science and technology love now. They loved them then too: well before its index listed Computers as a topic, Newsweek had entries under mathematical machines and electric brains.

The Machine Translation Fiasco

Machine translation is the most glaring casualty in the progress of AI: it was endorsed by reputable figures at major institutions, and was funded for more than a decade. It nevertheless subsequently proved impossible, given the science and technology available, by the late 1950s. The fact that the effort was premature in terms of the scientific knowledge and computer technology of the early 1950s, and that much research done under these auspices was entirely disengaged from complex information processing and artificial intelligence as intellectual movements, did not help. The poor public relations that machine translation caused reflected on these more viable endeavors, too.

Like physical automata, the goal of machine translation has an eternal quality to it. Reportedly, it was first conceived by Warren Weaver and A.D. Booth in 1946 (362). But a bit of contemplation turns up earlier sources. Esperanto, which was invented out of the stems of Western European languages in 1887, seems to anticipate mechanical translation. John Pfeiffer mentions several translating machines created by hobbyists as well (363). Booth tells us that in 1933 one Troyansky, “obtained a Russian patent for such a device [a mechanical dictionary]” (364).

In 1949, apparently elated with the success of the autonomous machinery of the Second World War and digital computing, mathematician Warren Weaver and engineer A.D. Booth, decided that it would be only a hop, skip, and a jump to simply automate translation from one language from another. Weaver stated, blithely:

“When I look at an article in Russian”, he wrote, “I say this is really written in English, but it has been coded in some strange symbols. I will now proceed to decode” (365).

“Translation”, his proposal for a grand project in automatic translation, was widely circulated (366). Easier said than done, it proved, but Weaver’s word was valued. His proposal that since we merely ‘decode’ from language to language, it must be pretty easily done artificially, was taken seriously. As mentioned in Chapter Two, Weaver was a distinguished electrical engineer who had been instrumental in the formation of information science. He had been manager of the National Defense Research Committee division, in which Norbert Wiener and company did early research in servomechanisms during the Second World War. Now he was in charge of natural science funding at the Rockefeller Foundation. In the late 1940s he remained at the peak of success. That same year, he and Claude Shannon published *The Mathematical Theory of Communication*.

This distinguished work did not mean that Weaver had any idea of how to proceed in MT. On the contrary, the latter book explicitly explains that its form of communication is about information minus the semantics. Emphasis on syntactic purity and minimal concern with semantics was fine for early computer languages, but not for linguistic nuances. That attitude should have kept Weaver away from Machine Translation- language is

fraught with confusing semantic details. His proposal was full of examples of the worst excesses of technological utopianism. Erroneous assumptions here demonstrate the hubris of information theory taken too far, reductionism in assuming that translation consists simply of simply replacing each word with equivalent words, and disregard for linguistics, syntax, idiom, semantic ambiguity, homonyms, and subtlety. Granted, Weaver apparently acquiesced to the need to create a kind of inter-linguistic Esperanto in between the input and the output languages, but he thought that these challenges would be easily surmounted (367).

Nor did these challenges deter the powers that be. The languages being automatically translated might well be English and Russian, so the military research establishment was interested enough to take the bait, and to continue to throw good money after bad for more than a decade. Automatic language processing acquired a journal, conferences, news reports, research groups both in the USA and abroad, and at least one thousand researchers during two decades after WWII (368). It attained prestigious addresses too. Engineer Anthony Oettinger, whose work was supported by the NSF and Rome Air Force Base, conducted his research at the Harvard Laboratory run by Howard Aiken (369). Yehoshua Bar-Hillel worked at MIT (370). With this sort of infrastructure and sunk costs, this rolling stone gathered considerable inertia as well.

The work itself, with the exception of the pioneering linguistics research by Victor H. Yngve, was in basis unsophisticated. Booth explains that the initial idea was simply storage of the stems of verbs in separate lookup tables for each language, the latter connected by adjacent machine addresses. This very description indicates just how early in the infancy of digital computing machine translation began. This look-up table with stems was augmented by separate sets of tables for declensions, etcetera.

The reported summit of this effort was a 'translation' that makes everyone who hears it wince. Reports of automatic translation program that turned "The spirit is willing but the flesh is weak" (English) into "The vodka is good but the meat is rotten" (Russian) began to circulate in the early 1960s. The oddity of this translation is not even the words, however. It is that it was what is

now called an urban legend. In 1995, MT researcher John Hutchins published a fascinating article in MT exploring the topic in detail and finding that this mock translation never actually took place. The phrase has been repeated so many times that it is not questioned (371).

Pfeiffer reports on some of the disasters of machine translation. For instance, the following sentences, taken from a speech by Nikita Khrushchev, are far better. The human English translation of the Russian original is:

“ While speaking about successes, we should always critically examine all aspects of our activity, we should not rest content with what has been achieved, we should constantly be concerned about complete utilization of the great reserves we have and of the possibilities for the powerful development of all branches of the national economy.”

The MT program turns the Russian speech into the following:

“ Talking about successes, we always should critical look at all side our activity, not calm on reached, constantly care about that in order to completely use having by us great reserves and possibility for high-power development all branch national economy.”(372)

Repeated reports of the “meat is rotten” translation fiasco and the like turned popular and scientific opinion against MT. Finally, in 1966, a National Academy of Sciences Committee report excoriated the field and its workers, resulting in the abrupt cancellation of many contracts (373). Funding sources practically fell off a cliff at this point. MT had much higher-quality results to show for itself, which work was ignored by the 1966 report.

Machine translation was given a bum rap: it was apparently more photogenic in its morbid fashion than the actual success of such programs as LT and GPS and timesharing, which spanned the same interval. It is important to refrain from poking fun at early machine translation, despite the dead end toward which the Oettinger work was obviously heading. This is because actually, machine translation was subject to very fruitful research even during the late 1950s as early as 1960 (374). Yngve’s COMIT computer language, first versions of which were presented in scientific circles in 1958, has led to further, unbroken progress in

machine translation (or linguistic data processing as it also known). Despite this, the field had been savaged by the press:

“ the coup de grace was a report prepared by a committee headed by John Pierce, then at Bell Labs, which concluded that a genuine machine translator was simply too ambitious at the moment, given the state of what was known about linguistics and about the use of natural language” (375).

However, there was a good deal of sincere but misguided scientific research, and garrulous talk to newsmagazine reporters, which was ill-conceived and destined to wind up in the pages of diatribes against AI. Some of this went by the name of AI, and while it is forgotten now, it garnered a good deal of publicity at the time. For instance, in 1960 Douglas Ross and Harrison Morse of MIT's Electronics Systems Laboratory composed a television script-writing language and program, SAGA II, for the Lincoln Laboratory TX-0. The program, operated in mnemonic code, defined several characters for a 'spaghetti western' script, along with illegal moves (robbers are not supposed to leave the scene without their loot, for instance). In October 1960, CBS broadcast the show, with live actors and a chunky script written by machine, under the title 'the Thinking Machine' (376).

Given such atrocities, perhaps part of the natural evolutionary process in which many start out but few survive, it is not surprising that AI was considered a science 'noir' field. The few giddy words that Newell and Simon happened to let slip were indeed taken seriously. Well-publicized quick and dirty faux Artificial Intelligence projects were as desirable as a third eye, but the scientists had to tolerate it. Serene as ever, Simon comments:

“ We [A. Newell and H.A. Simon] have adopted the policy (was it the anarchist Bakunin's ? or Sorel's ?) of propaganda of the deed, not propaganda of the word.” The best rhetoric comes from building and testing models and running experiments”. (377)

Chapter 13. Conclusion: Another Pregnant Pause

Achievements at the End of the Fifties

The achievements of the years between the late 1930s and the late 1950s were mammoth in number and magnitude.

As we observed in the first pages of this book, they include the introduction of a way to realize the idea of AI itself. The idea had been patently impossible during the entirety of prior human history. During the Forties, the digital computer was invented, as was information theory. Automata had been conceived and blessed by Turing and Von Neumann, and in the process inspired the new generation. The implementation of AI as an idea, even in the simple achievements of the Logic Theorist and in Minsky's Euclid proofs, was still an enormous advance. AI had lived a lingering existence as an idea that seemed as if it would always be in the future. This made the Dartmouth Conference both more revolutionary and more practical than might be thought.

In addition to the significant achievement of the incipience of AI itself, computing had advanced prodigiously during the years chronicled in this book. Each aspect of computing- input, output, storage, memory, programming languages- was under reiterated improvement and increasing demand far beyond the initial auspices of the military. The solely experimental ENIAC-EDVAC project had been barely functional and so inordinately demanding of electricity that its operation dimmed the lights of Philadelphia. The repeated and successive projects at various world facilities had improved the entire functionality and the specific components of digital computing.

These improvements were starting to revolutionize popular culture as well. In 1955 Sony introduced its TR 55 transistor radio, decades before the average person on the street would ever see a computer in person. The portable television would follow in 1959 (378). The advent of standardized programming conventions in the form of programming languages would replace 'coders' who wrote programs in machine language, with subsequently more abstract routines (written by the same coders, renamed 'programmers'). DEC computing was founded in 1957 (379). This MIT spinoff would sell computers which were far more amenable in input and output usage than were the leased IBM mainframes. The state of the art had indeed advanced well beyond giant brains on legs and anthropomorphic omniscient robots.

By the final years of the 1950s, AI as an engineering discipline that would iteratively inform and learn from cognitive psychology

was being developed by Newell and Simon. AI as a form of engineered artifacts was being developed by Minsky and McCarthy. McCarthy had already invented LISP. He had envisioned timesharing and was mustering colleagues to implement this programming convention. Each of the dozen or so AI practitioners effectively increased demand for improvements of every facet of computer practice.

Finally, the characteristic different strands of AI research were also established by the conclusion of the 1950s. It was quite obvious what differing discussions would take place repeatedly at the Thanksgiving dinner table- so to speak- for the next decades. John McCarthy would continue to work in list-processing languages, timesharing, and formalisms for the superior expression of semantics. Marvin Minsky, incessantly curious about every facet of AI, would develop engineered artifacts and theories of knowledge representation in iteration with artifacts. Allen Newell and Herbert Simon continued to work in cognitive psychology and in problem-solving in increasingly semantically complex fields. Each of these individuals would continue to enrich and enlarge their research, but none of them fundamentally altered. The ancestral or nodal concepts that each brought to that table did not change, but became successively more expressive with further scholarship. Given the great differences in these strands of research, the fact that from the beginning AI took place at several institutions- with Stanford added at the start of the 1960s- was tremendously fortuitous.

Another Pregnant Pause

One leitmotif of AI's earliest growth through the year of its establishment was its indigenous appearance from several sources- Rand, MIT, and the IAS and Princeton influence. Its founders would spend their lives as professors at major universities. Yet when the field was declared in 1956, only Herbert Simon was actually a tenured university professor, and only Newell and Simon's research group was well-funded and able to offer degrees and courses. As we saw in the chapter concerning Minsky and McCarthy, they achieved this several years later in the 1950s. But one university, or even two, were not enough a sufficient catalyst to propagate AI as a field of computer science.

McCarthy and Minsky, and their own students, would need research funding; timesharing itself integrally required both hardware and software development; input and output improvements would both require computer funding that squared, and cubed, all existing energy thrown into this field.

Notwithstanding brilliant research carried out by all of the parties, none of them could change the dynamics of research funding unilaterally. Thus, in the late 1950s, like the earlier pause in the postwar era, the field required external assistance to help it move forward. In the early 1950s, what had been needed was sheer human ingenuity, chance and fortuitous meetings, such as those between Newell and Selfridge; Newell, Shaw and Simon; Minsky and McCarthy, and the entire cast eventually forming a loose group in 1956. AI would not become a research program followed at many universities until computing was pulled into a central role in the United States military.

The Spark of Sputnik

AI needed to be needed- in fact, demanded- by a powerful benefactor. The best prospective rich uncle was the DoD, and this did indeed happen- deus ex machina, literally. In a brief outline, the magnitude of the events make themselves clear. On October 4, 1957, the Soviet Union sent Sputnik I, the first man-made earth satellite, into orbit around the Earth. It remained in orbit for exactly four months. The Soviets followed this foray with Sputnik II later in the month. The United States went ballistic, so to speak, and responded with accelerated development of its own projects, the Thor and Jupiter missiles. Eight more attempted satellite liftoffs initiated by the both parties, took place in 1958-1959 alone. Whatever the larger context of the ongoing arms race, the Sputnik missile launch permanently turned up the volume.(380)

By mid-January of 1958, President, and former General, Dwight Eisenhower had established a bureaucratic separation between the missile projects and other divisions within the Defense Department. This was apparently intended to prevent internal conflicts of interest within the Army and Air Force- both of which had missile projects- from impeding more dramatic progress on these missiles). To this end he also established the Advanced Research Projects Agency, to answer directly to the Secretary of

Defense, "for the unified direction and management of the antimissile missile program and for outer space projects".(381)

Over the next year, President Eisenhower took other actions which strengthened the scientific cohort within the military. He appointed James Killian, the President of MIT, as a presidential assistant for science, and Killian in turn immediately established the President's Science Advisory Committee. He also set in motion the creation of the position of Director of Defense Research and Engineering. This position superseded the existing one of Assistant Secretary of Defense for Research and Engineering, and provided further jobs within the military-intelligence-computing establishment. (Later ARPA executives, for instance George Heilmeyer, sometimes began their careers in DDR&E). Congress also began paying more attention to space sciences; the Senate established a Standing Committee on Aeronautical and Space Sciences.

ARPA had been placed in charge of missiles in specific and upper atmospheric exploration and outer space vehicles in general. However, it lasted in this role for less than one year. Its specific administration of missile projects was complemented and broadened by the foundation of the National Aeronautics and Space Administration. NASA itself replaced an earlier and less powerful agency. The National Advisory Committee for Aeronautics, a government body which carried out NASA's functions, already had existed since 1916. NACA was the proto-NASA, just as the ONR was the proto-ARPA. NACA was renamed and given \$125 million and a far more prominent position in 1958 (382).

Sputnik was manna from heaven to the United States' scientific community in general. The response was vastly increased spending on military research and development, especially where scientific education and experimentation were concerned. The United States was jolted into a frenzy of military and scientific spending:

" The nation responded by authorizing billions more in R&D-specifically, expenditures went from \$3 billion in 1957, the year of the Sputniks, to \$15 billion in 1964. Research aid to universities spurted upward during this period, as did budgets for research into all forms of education."(383)

But the foundation of ARPA, and of NASA in short order, was not immediately consequential for AI; the technical needs of aeronautics did not, at first, seem to include intelligent control. It was several years before this broader defense agenda saw a need to foster 'exotic' computer research. Between 1958 and 1962, ARPA functioned as an "interim space agency". By the early 1960s, it had settled on the goal of developing re-entry physics for missiles (384). Jack Ruina, ARPA Director from 1961 to 1963, stated that "The major programs we had were ballistic missile defense [and] nuclear test detection". Other things were on the margin of the larger defense picture (385).

The Unforeseeable Consequences of AI 1956, Increasingly Evident

These historical events of a global scale had nothing to do with AI intrinsically, but would bear great repercussions for the field. ARPA monies would permit and encourage the recruitment of dozens and later hundreds of researchers. When the Apollo program (i.e., NASA's Moon project) got underway several years later, it would bring about the immense improvement of the integrated circuit as well. This hardware improvement, combined with other measures that had more specific import for AI rather than generically for hardware for computing, would revolutionize the field itself and continues to do so.

But that is another story.

Chapter 14. Acknowledgements

Wandering through the University of California at Berkeley campus during graduate school, the author found herself in the back of an introductory lecture on Artificial Intelligence computing in Cory Hall. As is her manner, she took notes in the form of a cursory list, which turned into far more. A little list can be a dangerous thing. As we know from the history of AI and its predecessors, lists are often at the start of things- commandments, dictionaries, encyclopedias, Cyc, LISP and list processing languages. Learning more did not satiate but made me hungrier: definitions provoked more questions- understanding the terms being used, the development of central problems in AI,

the software projects being discussed, and their provenance and creators and chronology. The itch to know more could not be scratched away.

Building The Second Mind originated in the author's doctoral dissertation on the introduction of commercial expert systems in the 1980s- "Developmental Characteristics and Spatial Formation in the Commercialization of Knowledge Base System Shells, 1975-1991" (May 1993, Department of City and Regional Planning). As is typical for dissertations, if not for this author's intellectual compulsions, that work stayed close to its focus. This was the commercial introduction of expert systems to the turbulent world of software applications of the 1980s, their technological adaptation to that environment, and their mixed success in that endeavor. Stanford, CMU and MIT were considered insofar as their differing intellectual environments produced distinctive commercial cultures.

I did not wish to let go of this issue, and it did not want to let go of me either. In a fine example of emergent functionality, the story took on the proverbial life of its own, and its topics grew into the past and the present. The nucleus of the original thesis concerned the years between 1975 and 1991. In this book, that chronological period has been supplemented by another 'nuclear' narrative concerning the prehistory of AI as an idea- one of the most magical and impossible ideas in history- and the early development of AI parallel to the slow and tortuous improvement of the general-purpose digital computer. The transformation of metal into spirit is not, strictly speaking, magical but certainly seems to be when one tries to explain it afresh.

For a few years, the book had to give way to raising children. During the Naughts, I developed a nostalgic cargo cult of my own return to my book. In 2010, the children were finally more independent, one going so far as to move to New Hampshire to demonstrate this point. One Spring day, the author hopefully took up the project again. Returning to this existential task worth doing was immensely gratifying: There's no place like home.

A list of people who have contributed to this work, through conversations, interviews, or classes I took, includes:

Mie Augier, Avron Barr, Jennie Benton, Daniel Bobrow, Tim Bresnahan, Lee Brown, Bruce Buchanan, Manuel Castells, Marie DesJardins, Paul Edwards, the late Robert Engelmores, Adam Farquhar, E. A. Feigenbaum, Richard Fikes, Ken Ford, Peter Friedland, Doug Fridsma, Yolanda Gil, Ascuncion Gomez, Tom Gruber, Peter Hall, Peter Hart, Barbara Hayes-Roth, Rick Hayes-Roth, Torsten Heycke, Donna Holm, Yumi Iwasaki, Jennifer Joseph, Maggie Kimball, Craig Knoblock, Pat Langley, Roger Launius, Henry Lieberman, Henry Lowood, Roberto Mazzolini, the late John McCarthy, Robert E. McGinn, Bill Miller, Andy Neuschatz, H. Penny Nii, Nils Nilsson, Arthur Norberg, Peter Norvig, Pierre Omidyar, Judy O'Neill, Bertram Raphael, Raj Reddy, Bradley Rhodes, Nathan Rosenberg, Stuart Russell, Annalee Saxenian, the late Oliver Selfridge, Ted Shortliffe, the late Herbert A. Simon, Luc Steele, Mark Stefik, Lynn Andrea Stein, Marty Tennenbaum, Shirley Tessler, Peche Turner, Richard Waldinger, Richard Walker, Robert Wilensky, Terry Winograd, and Lofti Zadeh.

On a personal note, the author would like to thank Lucy Gill, Gary Hauser, Mary Heldman, Linda Herschenson, Paul Hufstedler, Gary Nagamori, the San Francisco Proust Reading Group, Maria Schopp, Mike Travers, Richard Waldinger, and Jennifer Yeh.

Finally, this book is dedicated to my daughters, Abigail Grace Skinner and Katherine Alice Skinner.

Chapter 15. Bibliography

Abbate, Janet. *Inventing the Internet*. Cambridge, MA: The MIT Press, 2000.

Abella, Alex. *Soldiers of Reason: the Rand Corporation and the Rise of the American Empire*. New York: Harcourt, 2008.

Aeronautics and Astronautics Chronology, compiled by NASA Historian Eugene M. Emme.
<http://www.hq.nasa.gov/office/pao/History/timeline.html>.

Agar, Jon. *The Government Machine: A Revolutionary History of the Computer*. Cambridge, MA: The MIT Press, 2003.

Akera, Atsushi and Nebeker, Frederik, Eds. *From 0 to 1: An authoritative history of modern computing*. New York, NY: Oxford University Press, 2002.

Akera, Atsushi. *Calculating a natural world: Scientists, Engineers and Computers during the rise of U.S. Cold War Research*. Cambridge, MA: The MIT Press, 2008.

Akera, Atsushi and Nebeker, Frederik, Eds. *From 0 to 1: An authoritative history of modern computing*. New York, NY: Oxford University Press, 2002.

Akera, Atsushi. *Calculating a natural world: Scientists, Engineers and Computers during the rise of U.S. Cold War Research*. Cambridge, MA: The MIT Press, 2008.

Armer, Paul. "Attitudes Toward Intelligent Machines". In Feigenbaum, Edward A. and Julian Feldman, eds. *Computers and Thought*. New York: McGraw-Hill. 1963: 389-405. Reprinted in Pylyshyn, Zenon, ed. *Perspectives on the Computer Revolution*. Englewood Cliffs, N.J. Prentice-Hall, Inc., 1970.

Ashby, Ross. "Design for a Brain". *Electrical Engineering*, December 1948, p379.

Ashby, W. Ross. "The Mechanism of Habituation". In Blake, D.V. and A.M. Uttley, eds. *Proceedings of the Symposium on the Mechanisation of Thought Processes*. National Physical Laboratories, Teddington, England. London: H.M. Stationary Office. 1959: 93-114.

Aspray, William, ed. *Computing Before Computers*. Ames, Iowa: Iowa State University Press, 1990.

Aspray, William. *John Von Neumann and the Origins of Modern Computing*. Cambridge, MA: The MIT Press, 1990.

Augier, Mie and James G. March, eds. *Models of a Man: Essays in Memory of Herbert A Simon*. Cambridge, MA: The MIT Press, 2004.

Babbage, Charles. *Charles Babbage: Passages from the life of a philosopher*. Edited, with a new introduction by Martin Campbell-Kelly. New Brunswick, NJ: Rutgers University Press, IEEE Press, 1994.

Backus, Jim." Can programming be liberated from the Von Neumann Style ? A functional style and its algebra of programs." *Communications of the ACM* 21, no. 8. (August, 1978): 613-641.

Backus, Jim." Programming in America in the 1950s: Some Personal Impressions". In Metropolis, N. et al. *A History of Computing in the Twentieth Century*. New York: Harcourt Brace Jovanovich. 1980: 125-135.

Bar-Hillel, Yehoshua." The present status of automatic translation of languages". *Advances in computers*. New York: Academic Press, 1960.

Barr, Cohen and Edward A. Feigenbaum. *The Handbook of Artificial Intelligence*. Volumes 1-3. Los Altos, CA: William Kaufmann, 1981.

Barr, Avron, Paul R. Cohen, and Edward A. Feigenbaum. *The Handbook of Artificial Intelligence*. Volume IV. Menlo Park, CA: Addison-Wesley Publishing Company, 1989.

Baum, Claude. *The System Builders: the story of SDC*. Santa Monica, CA: The System Development Corporation, 1981.

Baumgartner Peter, and Payr, Sabine, eds. *Speaking Minds: Interviews with Twenty Eminent Cognitive Scientists*. Princeton, NJ: Princeton University Press, 1995.

Benford, Gregory and Elisabeth Malartre. *Beyond Human: Living with Robots and Cyborgs*. New York: A Forge Book, 2007.

Bennett, Edward; James Degan, and Joseph Spiegel, eds. *Military Information systems: the design of computer aided systems for command*. New York: Praeger, 1964.

Berkeley, Edmund C. *Giant brains; or, Machines that Think*. New York: Wiley, 1949.

Berners-Lee, Tim, with Mark Fischetti. *Weaving the Web: The Original Design and ultimate destiny of the World Wide Web*. New York: HarperBusiness, 2000.

Bernstein, Jeremy. *Science observed: essays out of my mind*. New York: Basic Books, 1982.

Birkhoff, G." Computing Developments 1935-1955, as Seen from Cambridge, USA". In Metropolis, N. et al. *A History of Computing in the Twentieth Century*. New York: Harcourt Brace Jovanovich. 1980: 21-30.

Blackburn, *Oxford Dictionary of Philosophy*. New York: Oxford University Press, 1994.

Blake and Uttley. In Blake, D.V. and A.M. Uttley, eds. *Proceedings of the Symposium on the Mechanisation of Thought Processes*. National Physical Laboratories, Teddington, England. London: H.M. Stationary Office. 1959.

Booth, A.D. *Digital Computers in Action*. New York: Pergamon Press, 1965.

Boring, E." Mind and Mechanism," *American J. of Psychology* 2, April 1946.

Bowden, B.V., ed. *Faster than thought*. London: Pitman Publishing, 1953.

Bowker, Geoff. "How to be Universal: Some Cybernetic Strategies". *Social Studies of Science* SAGE Publishers, Newbury Park, N.J., 23 (1993): 107-127.

Bradbury, Ray. *Fahrenheit 451*. New York: Ballantine Books (DelRey), 1953.

Brendon, Piers. *The Dark Valley: A panorama of the 1930s*. New York: Vintage, 2000.

Bromley, Allan. G. "Analog Computing Devices". In *Computing Before Computers*, edited by William Aspray, 156-199. Ames, Iowa: Iowa State University Press, 1990.

Bromley, Allan. G. "Difference and Analytical Engines". In W. Aspray, ed. *Computing Before Computers*. Ames, Iowa: Iowa State University Press. 1990: 59-98.

Bruner, Jerome S.; Goodnow, Jacqueline; and George A. Austin. *A Study of Thinking*. New York: John Wiley & Sons. 1956.

Bush, Vannevar. "As We May Think", *The Atlantic Monthly*. Summer 1945.

Bush, Vannevar. *Science, the Endless Frontier*. 1945. Reissued as part of the Tenth Anniversary Observance National Science Foundation 1950 1960. Washington D.C.: National Science Foundation, 1960.

Campbell-Kelly, Martin and William Aspray. *Computer: a history of the information machine*. 1st edition. New York: Basic Books, 1996.

Campbell-Kelly, Martin and William Aspray. Computer: a history of the information machine. Second edition. Boulder, CO: Westview Press. 2000.

Campbell-Kelly Martin. From Airline Reservations to Sonic the Hedgehog: A History of the Software Industry. Cambridge, MA: The MIT Press, 2004.

Campbell-Kelly, Martin." Punched-Card Machinery". In W. Aspray, ed. Computing Before Computers. Ames, Iowa: Iowa State University Press. 1990: 122-155.

Capek, Karel. R.U.R. (Rossum's Universal Robots). A play in three acts and an epilogue. Translated from the Czech by P Selver and adapted for the English stage by Nigel Playfair. London and New York: Oxford University Press. 1961. Originally published 1923.

Capek, Karel. The Insect Play. London and New York: Oxford University Press. 1961. Originally published 1923.

Ceruzzi, Paul E." Electronic Calculators". In W. Aspray, ed. Computing Before Computers. Ames, Iowa: Iowa State University Press. 1990: 223-250.

Ceruzzi, Paul E." Relay Calculators". In W. Aspray, ed. Computing Before Computers. Ames, Iowa: Iowa State University Press. 1990: 200-222.

Ceruzzi, Paul E. A History of Modern Computing: Second Edition. Cambridge, MA: The MIT Press, 2003.

Charniak, Eugene and Drew McDermott. Introduction to Artificial Intelligence. Reading, MA: Addison-Wesley, 1985.

Clynes, Manfred, and Nathan S Kline, " Cyborgs and Space," Astronautics September 1960.

Conant, Jennet. Tuxedo Park: A Wall Street Tycoon and the Secret Palace of Science that changed the course of World War Two. New York: Simon & Schuster. 2003.

Conway, Flo and Jim Siegelman. Dark hero of the Information Age: In search of Norbert Wiener, the father of Cybernetics. New York: Basic Books: 2005.

Copeland, B. Jack, ed., The Essential Turing: The ideas that gave birth to the Computer Age. Oxford: Clarendon Press, 2004.

Cote, Alfred J. The Search for the Robots. New York: Basic Books, 1967.

Craik, Kenneth. The Nature of Explanation. 1943.

Crevier, Daniel. AI: The Tumultuous History of the Search for Artificial Intelligence. New York: Basic Books, 1993.

Crowther-Heyck, Hunter. Herbert A. Simon: The Bounds of reason in Modern America. Baltimore: The Johns Hopkins University Press, 2005.

Davis, M.D." A note on Universal Turing Machines", in Shannon, C. E., and J. McCarthy, eds. Automata studies. Princeton: Princeton University Press. 1956: 167-176.

Denicoff, M." AI Development and the Office of Naval Research". In Bartee, T., ed. Expert Systems and Artificial Intelligence: Applications and Management. Philadelphia: Howard W. Sams & Company (MacMillan). 1988: 271-289.

Dickson, Paul. Think Tanks. New York: Atheneum, 1971.

Dreyfus, H.L. What Computers Can't Do: The Limits of Artificial Intelligence. Second edition. New York, N.Y.: Harper & Row, 1979.

Eatwell, Milgate, and Newman, eds. Allocation, Information, and Markets. The New Palgrave Dictionary of Economics (1990).

Eden, Murray. "Cybernetics" article in Machlup and Mansfield 1983.

Edwards, Paul N. The closed world: computers and the politics of discourse in Cold War America. Cambridge, MA: MIT Press, 1996.

Engelmore, Robert S. "AI Developments: DARPA and ONR Viewpoints". In Bartee, T., ed. Expert Systems and Artificial Intelligence: Applications and Management. Philadelphia: Howard W. Sams & Company (MacMillan). 1988: 213-218.

Etkowitz, Henry. MIT and the rise of entrepreneurial Science. New York: Routledge, 2002.

Evans, Christopher. The Making of the Micro: a History of the Computer. London: Victor Gollancz, Ltd., 1981.

Evans, Christopher. The Micro Millennium. New York: Washington Square Press, 1979.

Feigenbaum, Edward A. and Julian Feldman, eds. Computers and Thought. New York: McGraw-Hill, 1963.

Feigenbaum, E.A., "The Simulation of Verbal Learning Behavior". In Computers and Thought.

Flamm, Kenneth. Creating the Computer: Government, Industry and High Technology. Washington, D.C.: The Brookings Institution, 1988.

Flamm, Kenneth. Targeting the Computer: Government Support and International Competition.

Foerstel, Herbert N. Secret science: Federal control of American science and technology. Westport, CT: Praeger, 1993.

Ford, Kenneth M., Clark Glymour and Patrick J. Hayes, eds. Android epistemology. Menlo Park, CA: AAI Press, 1995.

Franchi, Stefano and Guven Guzeldere, eds. Mechanical Bodies, Computational Minds: Artificial Intelligence from automata to cyborgs. Cambridge, MA: The MIT Press, 2006.

Frenkel, James, ed. True Names By Vernor Vinge and the opening of the cyberspace frontier. New York, N.Y.: TOR, 2001.

Goldstine, Herman H. The Computer from Pascal to Von Neumann. Princeton, N.J.: Princeton University Press, 1972.

Gottfried, T. Alan Turing: The Architect of the Computer Age. Danbury, CT: Franklin Watts, 1996.

Grosch, Herbert R. Computer: bit slices from a life. Novato, California: Underwood-Miller, 1991.

Hapgood, Fred. Up the infinite corridor: MIT and the technical imagination. Reading, MA: Addison-Wesley Pub. Co., 1993.

Heims, Steve J. John von Neumann and Norbert Wiener: From Mathematics to the Technologies of Life and Death. Cambridge, MA: The MIT Press, 1980.

Heims, Steve J. The Cybernetics Group. 1991.

Hilts, Philip J. Scientific Temperaments: Three Lives in Contemporary Science. New York: Simon & Schuster, 1982.

Hiltzik, Michael. Dealers of lightning. Xerox Parc and the dawn of the computer Age. New York: Collins Business, 1999.

Hobbes, Thomas. Body, man, and citizen. Selections. 1656. Edited with an introduction by Richard S. Peters. New York: Collier Books. 1962.

Hodges, A. Turing: the Enigma of Intelligence.

Hutchins, John. "The Whiskey was invisible', or Persistent myths of MT," MT News International (June 1995): 17-18.

Inglis, Fred. *The Cruel Peace: Everyday Life and the Cold War*. New York: BasicBooks, 1991.

James, William. "The sentiment of rationality". *Writings, 1878-1899*. New York: Library Classics of the United States. The Library of America, 1992.

Jeffress, L.A., ed. *Cerebral Mechanisms in Behavior: The Hixon Symposium*. New York: John Wiley, 1951.

Kaplan, Fred. *The Wizards of Armageddon*. New York: Simon & Schuster, 1983.

Killian, Jr., James R. *Sputnik, Scientists, and Eisenhower: A Memoir of the First Special Assistant to the President for Science and Technology*. Cambridge, MA: The MIT Press, 1977.

Kuhn, Thomas S. *The Structure of Scientific Revolutions*. Second Edition. Chicago: University of Chicago Press, 1970.

Lakatos, Imre. John Worrall and G. Currie, eds. *The methodology of scientific research programmes*. Cambridge, U.K.: Cambridge University Press, 1978.

Lashley, K.S. "The Problem of Serial Order in Behavior," In L.A. Jeffress, ed. Cerebral Mechanisms in Behavior: The Hixon Symposium. New York: John Wiley. 1951: 112-136.

LeCuyer, Christophe. *Making Silicon Valley Innovation and the growth of high tech*. Cambridge, MA: The MIT Press, 2006.

Wilhelm Gottfried Leibniz. *Monadology*.

Leslie, Stuart. *The Cold War and American Science: The Military-Industrial-Academic Complex at MIT and Stanford*. New York: Columbia University Press, 1993.

J.C.R. Licklider, Charles Babbage Institute Oral History Interview 150, 1988.

Lucky, Robert W. *Silicon Dreams: information, man and machine*. New York: St. Martins Press, 1989.

Machlup, Fritz, and Una Mansfield, eds. *The Study of Information: Interdisciplinary Messages*. New York: John Wiley & Sons, 1982.

Markoff, John. *What the dormouse said: How the Sixties counterculture shaped the personal computer industry*. New York: Penguin Books, 2005.

Mayr, Otto. *The Origins of Feedback Control*. 1970.

Mazlish, Bruce. "The Man-Machine and Artificial Intelligence", in Franchi and Guzeldere eds., 2006.

McCarthy, John. "Programs with Common Sense". In Blake, D.V. and A.M. Uttley, eds. Proceedings of the Symposium on the Mechanisation of Thought Processes. National Physical Laboratories, Teddington, England. London: H.M. Stationary Office. 1959: 75-84.

McCarthy, John. Interview by William Aspray, audio tape(s) and transcript, Palo Alto, Calif., 2 March 1989, Charles Babbage Institute Interview. OH 156.

McCarthy, John. "The Inversion of Functions Defined by Turing Machines," In Shannon, C. and J. McCarthy, eds. "Automata Studies", Annals of Mathematical Studies N34, Princeton, N.J.: Princeton University Press. 1956: 177-182.

McClintock, Robert. "Machines and Vitalists: reflections on the ideology of cybernetics", The American Scholar. Spring 1966: 249-255.

McCorduck, Pamela. *Machines who think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. San Francisco, CA: W.H. Freeman and Company, 1979.

McCulloch, Warren Sturgis. "Why the Mind is in the Head". In L.A. Jeffress, ed. *Cerebral Mechanisms in Behavior: The Hixon Symposium*. New York: John Wiley. 1951: 42-57. Transcript of discussion attached to text.

McCulloch, Warren. "Agatha Tyche: of nervous nets - the lucky reckoners." In Blake, D.V. and A.M. Uttley, eds. *Proceedings of the Symposium on the Mechanisation of Thought Processes*. National Physical Laboratories, Teddington, England. London: H.M. Stationary Office. 1959: 611-626.

McDougall, Walter A. *The Heavens and the Earth: A Political History of the Space Age*. New York: Basic Books, 1985.

Michie, Donald and Rory Johnston. *The knowledge machine: artificial intelligence and the future of man*. New York: W. Morrow, 1985.

Minsky, Marvin Lee. "Some methods of Artificial Intelligence and heuristic programming". In Blake, D.V. and A.M. Uttley, eds. *Proceedings of the Symposium on the Mechanisation of Thought Processes*. National Physical Laboratories, Teddington, England. London: H.M. Stationary Office. 1959: 3-28.

Minsky, Marvin Lee. Interview by Arthur L. Norberg, audio tape(s) and transcript, Cambridge, Mass., 1 November 1989, Charles Babbage Institute Interview. OH 179.

Minsky, Marvin, and Papert, Seymour. *Perceptrons: an introduction to computational geometry*. Cambridge, MA: The MIT Press, 1969. Expanded edition 1988.

Morrison, Philip and Emily Morrison, eds. *Charles Babbage and his calculating engines; selected writings by Charles Babbage and others*. New York, Dover Publications, 1961.

Morrison, Philip and Emily Morrison, eds. Charles Babbage on the principles and development of the calculator. and other seminal writings by Charles Babbage and others. New York, Dover Publications, 1961.

de Menabrea, Luigi. Sketch of the Analytical Engine Invented by Charles Babbage, Esq. Bibliotheque Universalle de Geneve. No.82, October, 1842. Translated with editorial notes by Augusta Ada, Countess of Lovelace. This work is reprinted in the Morrisons' volume.

Nemes, T.N. Cybernetic Machines. Trans. from the Hungarian by I. Foldes. New York: Gordon and Breach, Scientific Publishers, Inc., 1962.

Newell, A." The Chess Machine: an Example of Dealing with a Complex Task by Adaptation." Proceedings of the 1955 Western Joint Computer Conference. Institute of Radio Engineers, NY: 101-107.

Newell, Allen. Interview by Arthur L. Norberg, audio tape(s) and transcript, Pittsburgh, Pa., 10-12 June 1991, Charles Babbage Institute Interview. OH 227.

Newell, Allen. BIOGRAPHICAL MEMOIRS. National Academy of Sciences. March 19, 1927 — July 19, 1992. By Herbert A. Simon.

Newell, A., J.C. Shaw, and H.A. Simon." Chess-Playing Programs and the Problem of Complexity". IBM Journal of Research and Development 2, 4. (October, 1958): 320-335. Reprinted in Computers and Thought.

Newell, Allen. Unified Theories of Cognition. Cambridge, MA: Harvard University Press, 1992.

Newell and Simon 1972. Human Problem Solving.

Norberg, Arthur L. and Judy E. O'Neill. Transforming computer technology: information processing for the Pentagon, 1962-1986. With contributions by Kerry J. Freedman. Johns Hopkins studies in the history of technology. Baltimore: Johns Hopkins University Press, 1996.

Oettinger, Anthony G. Automatic Language Translation: Lexical and technical aspects with particular reference to Russian. Cambridge, MA: Harvard University Press, 1960.

Ornstein, Severo. Computing in the middle ages: A view from the trenches, 1955-1983. 1st books library. 2002.

Ornstein, Severo. Interview by Judy E. O'Neill, audio tape(s) and transcript, Woodside, Calif., 6 March 1990, Charles Babbage Institute Interview. OH 183.

Papert, Seymour, in McCulloch 1965.

Parrish, Thomas. The Cold War Encyclopedia. A Henry Holt Reference Book. New York: Henry Holt and Company, 1996.

Penick, James L. Jr.; Carroll W. Pursell, Jr.; Morgan B. Sherwood; and Donald C. Swain, eds. The Politics of American science, 1939 to the Present. Revised edition. Cambridge, MA: The MIT Press, 1972.

Peterson, Institute Historian T. F. Nightwork: A History of Hacks and Pranks at MIT. Cambridge, MA: The MIT Press, 2003.

Petzold, Charles. The Annotated Turing: A guided tour through Alan Turing's historic paper on Computability and the Turing Machine. Indianapolis, IN: Wiley, 2008.

Pfeiffer, John. The thinking machine: everyman's introduction to the world of electronic devices. Philadelphia: Lippincott, 1962.

Pitts, W. and Miller, G. "Comments on Session on Learning Machines". Proceedings of the 1955 Western Joint Computer Conference. Institute of Radio Engineers, NY: 108-110.

Polya, G. How to solve it. Second edition. Garden City, NY: Doubleday Anchor, 1957.

Pylyshyn, Zenon, ed. Perspectives on the Computer Revolution. Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1970.

The RAND Corporation: the First Fifteen Years. Santa Monica, CA, 1963.

Raphael, Bertram. The Thinking Computer: Mind Inside Matter. San Francisco, CA: W.H. Freeman & Co., 1976.

Redmond, Kent C., and Thomas M. Smith. Project Whirlwind: the History of a Pioneer Computer. Bedford, MA: Digital Press, 1980.

Redmond, Kent C., and Thomas M. Smith. From Whirlwind to Mitre: the R&D Story of the SAGE Air Defense Computer. Cambridge, MA: The MIT Press, 2000.

Rheingold, Howard. Tools for Thought: The People and Ideas behind the Next Computer Revolution. New York: Simon and Schuster, 1985.

Rhodes, Richard. The Making of the Atomic Bomb. New York: Simon & Schuster, 1986.

Reid, T. R. The Chip: how two Americans invented the microchip and launched a revolution. New York: Simon and Schuster, 1984.

Rice, Elmer L. The Adding Machine. Samuel French, New York, NY., 1923, 1956.

Rosenblatt, Frank." The Perceptron: A Probabilistic model for information storage and organization in the brain," Psychological Review LXV (1958): 386-407.

Rosenblatt, Frank." Perceptron Experiments," Proceedings of the IRE XLVIII (1960): 301-309.

Rosenblueth, Wiener, and Bigelow." Behavior, Purpose, and Teleology". Philosophy of Science (January 1943): 18-24.

Saxenian, AnnaLee. Regional Advantage: Culture and Competition in Silicon Valley and Route 128. Cambridge, MA: Harvard University Press, 1994.

Shannon, C.E., "Problem-solving electric mouse aids in improved Telephone Research," Electrical Engineering, July 1952, p671.

Shannon, Claude. Interview. FF59. CMU archives. Pamela McCorduck Collection. Transcripts, research, related materials. W.W. Bledsoe Interview. FF59.

Shannon, Claude." Computers and automata", In Pylyshyn, Zenon. ed. Perspectives on the Computer Revolution. Englewood Cliffs, N.J.: Prentice-Hall. 1970: 114-127.

Shannon, C. and Weaver, W. The mathematical theory of communication. 1963. Urbana: University of Illinois Press. (Originally 1949).

Shannon, C. E." A Universal Turing Machine with two internal states". In Shannon, C. E., and J. McCarthy, eds. Automata studies. Princeton: Princeton University Press. 1956: 157-166.

Shannon, C. E., and J. McCarthy, eds. Automata studies. Princeton: Princeton University Press. 1956.

Shannon, Claude. "A Chess-Playing Machine". Scientific American. February 1950: 48-51.

Shasha, Dennis, and Cathy Lazere. Out of their minds: The Lives and discoveries of 15 great computer scientists. New York: Copernicus, An Imprint of Springer-Verlag, 1998.

Shurkin, Joel N. Engines of the mind: the Evolution of the computer from mainframes to microprocessors. 1984. New York: Norton. Updated paperback edition 1996.

Simon, Herbert. Models of bounded rationality. Volume Three. Cambridge, MA: The MIT Press, 1997.

Simon, Herbert A. Models of man: social and rational; mathematical essays on rational human behavior in a social setting. New York, Wiley, 1957.

Simon, Herbert. The Sciences of the Artificial. Cambridge, MA: The MIT Press. Second Edition, 1981.

Simon, Herbert A. Models of Thought. New Haven: Yale University Press, 1979.

Simon, H.A. and A. Newell. "Heuristic Problem Solving: The Next Advance in Operations Research." Operations Research 6 (January-February 1958): 1-10.

Simon, Herbert A.; Donald W. Smithburg, and Victor A. Thompson. First edition, New York, Knopf, 1950. Administrative behavior by Herbert A. Simon, (New York, Macmillan Co. 1947.

Simon, H.A. "The behavioral theory of the firm", 1963.

Simon, H.A. "Rational choice and the structure of the environment", 1956.

Simon, Herbert Alexander. Models of my Life. 1991.

Simon, H.A. and A. Newell. "Heuristic Problem Solving: The Next Advance in Operations Research". Operations Research 6 (January-February 1958): 1-10.

Skinner, Rebecca Elizabeth. "Developmental Characteristics and Spatial Formation in the Commercialization of Knowledge Base System Shells, 1975-1991". Department of City and Regional Planning, University of California at Berkeley, May 1993.

Slocum, Jonathan. Machine Translation Systems (studies in natural language processing. New York: Cambridge University Press, 1987.

Snow, C. P. The Two Cultures and a Second Look. An Expanded Version of The Two Cultures and the Scientific Revolution. 1959. Cambridge, U.K.: The University Press, 1978.

Standage, Tom. The Victorian Internet: the remarkable story of the telegraph and nineteenth-century's On-Line Pioneer. New York: Walker & Company, 1998, 2007.

Standage, Tom. The Turk: The life and times of the famous eighteenth-century chess-playing machine. New York: Walker & Company, 2002.

Taube, Mortimer. Computers and common sense; the myth of thinking machines. New York: Columbia University Press, 1961.

Taylor, Rattray. "The Age of the androids". from Encounter, November 1963; In Pylyshyn, Zenon, ed. Perspectives on the Computer Revolution. Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1970.

Turing, A.M. "Computing Machinery and Intelligence". (1950). In Feigenbaum, Edward A. and Julian Feldman, eds. Computers and Thought. New York: McGraw-Hill. 1963: 11-39.

Turing, A.M., "Intelligent Machinery". (1947). Machine Intelligence 5. B. Meltzer and D. Michie, eds. Edinburgh: Edinburgh University Press. 1969.

Turing, Alan." The Chemical Basis of Morphogenesis", 1952, in B. Jack Copeland, ed., The Essential Turing: The ideas that gave birth to the Computer age. Oxford: Clarendon Press. 2004: p519-561.

Turkle, Sherry. The Second Self.

Turner, Fred. From Counterculture to Cyberculture: Stewart Brand, the whole earth network, and the rise of digital utopianism. University of Chicago Press, 2006.

Von Neumann, John. The Computer and the Brain. 1957. Second Edition, with a Foreword by Paul M. Churchland and Patricia S. Churchland. New Haven, CT: Yale University Press, 2000.

Waldrop, M. Mitchell. The Dream Machine: J.C.R. Licklider and the revolution that Made Computing Personal. New York: Viking, 2001.

Walker, Martin. The Cold War: A History. New York: Henry Holt and Company, 1993.

Walter, Grey." An Imitation of life". Scientific American May, 1950.

Walter, Grey." Totems, toys and tools,". Reprinted in Pylyshyn, Zenon, ed. Perspectives on the Computer Revolution. Englewood Cliffs, N.J. [1953]; 1970.

Watson's 1913 Lectures at Columbia University Psychological Seminary; printed in the Psychological Review 20: 158-177.

Weaver, Warren." Translation". In Locke, William Nash and A. Donald Booth, eds. Machine translation of languages; fourteen essays. [Cambridge] Published jointly by Technology Press of the Massachusetts Institute of Technology and Wiley, New York, 1955.

Wiener, Norbert. "Some moral and technical consequences of automation", Science (May 6, 1960): 1355-1359.

Wiener, Norbert. *The Human Use of Human Beings: Cybernetics and Society*. Boston: Houghton Mifflin Company, 1950.

Williams, Michael R. *A History of Computing Technology*. Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1979.

Winston CBI OH interview, 1990.

Wood, Gaby. Edison's Eve: A Magical History of the Quest for Mechanical Life. New York: Anchor Books. 2002.

Wooldridge, Dean E. *The machinery of the brain*. 1963.

Wooldridge, Dean E. *Mechanical Man: the Physical Basis of Intelligent Life*. New York: McGraw-Hill, 1968.

Zachary, G. Pascal. *Endless Frontier: Vannevar Bush Engineer of the American Century*. Cambridge, MA: The MIT Press, 1999.

Chapter 16. Endnotes

(1) Nathan Rosenberg, *Inside the Black Box: Technology and Economics*, 1982.

(2) The literature on general-purpose technologies has appeared during the past two decades. E. Helpman's edited volume is the best overall introduction.

(3) References to include Paul Edwards' *The Closed World*, Mirowski, Licklider biography, Jon Agar's *The Government Machine*.

(4) This is an affectively appealing and half-true thesis. *Hackers: Heroes of the Computer revolution*, and *Where wizards stay up*

late, to name two of the most appealing such works, argue that revolutionary hackers created most if not all of the astonishing inventions in computer applications. The best immediate refutation of this theory- at least as a general statement- is that the hackers further developed computing after it was already well-initiated, or at least had been in development for slightly under two decades. Anyone who thinks that putative participation in a counterculture can be defined by one's style of dress should ponder the examples of Vladimir Lenin and Malcolm X, both of whom always wore a dress shirt with a tie.

(5) Most scholarship historical work on the history of the computer credits the government, specific individuals and organizational cultures with the advances of the computer and the Internet. Several examples of this historical literature include:
Hafner and Lyon, *Where Wizards Stay Up Late: The Origins of the Internet*, 1998.
Hiltzik, *Dealers of lightning. Xerox Parc and the dawn of the Computer Age*, 1999.
Hafner and Markoff, *Cyberpunk: Outlaws and Hackers on the Computer Frontier*, 1992.

(6) Histories of computing and of the regions (Silicon Valley, Route 128) it was fostered tend to emphasize the cultural and organizational environment. This is a considerable literature, including the works of Atsushi Akeru, Freiburger and Swaine, Lowood, Norberg and O'Neill, O'Mara, Saxenian, and Waldrop among others.

(7) Paul Thagard's classic work, Thagard, Paul R. Conceptual Revolutions, 1992, comes to mind. The adoption of AI as a central engineering and scientific problem to be undertaken entailed a paradigm transition in which the characteristic problems to which people devoted their attention itself changed.

(8) Paul Edwards' *The Closed World* evokes the total environment of the Cold War research environment, as do the works of Stuart Leslie and P. Mirowski, among others. However, Western literary elite culture of the time was similarly convinced of its own correct

interpretation of history, and similarly culturally powerful. Two 'Closed Worlds' don't create an open world. The end of the Cold War, and the appearance of the nearly universally accessible internet, would truly end the closure of both environments. AI, as part of the vanguard of computing, would contribute to the opening of both closed worlds in the 1980s and beyond.

(9) J. McCarthy; M.L. Minsky, N. Rochester, and C.E. Shannon." A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence".

(10) Ceruzzi cites the oft-quoted prediction, made by Howard Aiken of Harvard University; A History of Modern Computing, p13.

(11) Ceruzzi's History of Modern Computing discusses this narrative in authoritative detail.

(12) William James, The sentiment of rationality". Writings, 1878-1899. 1992, 534.

(13) "Heuristic or 'ars inveniendi', was the name of a certain branch of study not very clearly circumscribed, belonging to logic, or to philosophy, or to psychology, often outlined, seldom presented in detail, and as good as forgotten today. The aim of heuristic is to study the methods and rules of discovery and invention. A few traces of such study may be found in the commentators of Euclid; a passage of Pappus is particularly interesting in this respect. The most famous attempts to build up a system of heuristic are due to Descartes and to Leibniz..."
Polya, How to Solve It, 113.

(14) Euclid's proof for the area of a triangle was later the base for one of the earliest formal searches, that is, Marvin Minsky's work presented at the 1956 Dartmouth Conference.

(15) Heron also invented the formula for determining the area of a triangle.

(16) Raphael, *The Thinking Computer*, 253. The word derives from a Hebrew term for 'embryo' or anything incompletely developed'. The myth itself is intriguing, as it mixes two divergent cultural strains in Judaism. The Golem myth partakes of the Jewish Counter-Enlightenment, a ex-post medieval stew of numerology, charismatic leadership, incessant prayer, and the cultivation of ecstatic states. The latter movement led to the formation of Hasidic sects in the depths of rural Eastern Europe in the eighteenth century. At the same time, the Mitnagdim, a reference in Hebrew to "those who are against", rejected mythology and charisma and endorsed scholarship and legalistic thinking. Thus, the apparently fantastical scientific advance of the Golem is magical rather than technological. Alchemy was "a medieval chemical science and speculative philosophy aiming to achieve the transmutation of the base metals into gold, the discovery of a universal cure for disease, and the discovery of a means for indefinitely prolonging life." (Webster's Ninth New Collegiate Dictionary). The science's goals were as elusive as the creation of the Golem.

(17) Aspray *Computing before Computers*, p104-106. Lull: Glymour, Ford, and Hayes (1995, 5) further report that Lull had hoped to use this device for missionary efforts: "...Lull believed that infidels could be converted if they could be brought to see the combinations of God's attributes...he thought that a representation of those combinations could be effectively presented by means of appropriate machines and that supposition was the key to his new method." The Columbia Encyclopedia reports that these efforts to convert the 'infidels' resulted in Lull's repeated deportation from Tunisia, and finally to his being stoned to death.

(18) Morrison and Morrison eds., 1961, xxiii.

(19) Simon and Newell, *Operations Research* 1958.

(20) Punched cards were also decisive for the industrial revolution: the Jacquard loom had further effects on other industries such as heavy construction and other work in textiles.

See Cardwell 186-189; Usher 1954, p289-295; Mokyr 1990 p101-103; Rosenberg and Vincenti 1978, p39.

(21) Morrison and Morrison 1961, and Swade 1991.

(22) Rationalist philosophy, originating with Descartes, epitomizes the mechanical view of the material world- and the human body and mind- as things that are explicable and knowable. Gaby Wood, *Edison's Eve: A Magical History of the Quest for Mechanical Life*, 2002, covers this topic in extensive and fascinating detail. Automata are discussed in Aspray 1990; Standage, *The Turk: The life and times of the famous eighteenth-century chess-playing machine*. Mazlish 2006.

(23) To say that this is a topic worthy of further elaboration is an understatement. For the treatment that the history of the computer deserves, the reader should pursue the following: Agar, *The Government Machine* 2003; Aker and Nebeker, Eds. *From 0 to 1: An authoritative history of modern computing*, 2002. Aspray, ed. *Computing Before Computers*, 1990; Campbell-Kelly and Aspray, *Computer: a history of the information machine*, 1996; Ceruzzi, *A History of Modern Computing: Second Edition*, 2003; and Cortada, ed., *Before the Computer*.

(24) Before the late 1940s, automata were typically mechanical rather than electrical, and emulated physiology. Increasingly in the 20th century they emulated function rather than the form of living things. See Walter 1953, p184; Wood 2002; Benford and Malartre 2007; and Bruce Mazlish 2005 also presents the history of automata.

The 'Lux Protozoon', created by the German F. Lux, was an ingenious control device:

" This device is a model of a unicellular animal. It consists of a short ribbed cylinder with two slots. According to the idea of its inventor, it is anchored in a brook, half submerged in the water. The water in the protozoon is considered as the food it is in the process of digesting. If there is too much water, it is "overfed"

then the rubber cover distends and the lever... closes the lower contact... and the small flap restricts the slot called "mouth".
Nemes, *Cybernetic Machines*, 164.

Control of input and output- both of which were water, in either case- was exerted in an analogue method in this artifact. The rubber cover reacted directly to the water, rather than reacting indirectly through the mediation of its nervous system. Various automata emulated vision as well as form, resulting in a mechanical dog, or dogs. The Hammond mechanical dog (1915) and the Philips watch dog, built during the 1920s, had photoelectric cells for eyes. When light shone on one eye, the motor turned on and moved the dog toward the light (light triggered the electric reaction, hence the label photoelectric). The first dog followed a moving light, while the second 'watch dog' stopped when the light grew too bright, and began howling. Pylyshyn, ed., 1970, p160; Nemes 1962, p164).

In AI proper, theories which explain a given phenomenon by means of a black box, rather than explaining it, are called 'homunculus' theories.

(26) Berkeley 1949, p7.

(27) Eden, 423.

(28) Freud And The Americans.

(29) Finally, Phenomenological psychology, a mostly European Continental tradition that described consciousness and mental states in an almost artistic fashion, did not become influential in the United States until the 1960s.

(30) Definition of Behaviorism, *The Mind's new science*.

(31) From Watson's 1913 Lectures at Columbia University Psychological Seminary; printed in the *Psychological Review* 20: 158-177.

(32) For the hundredth time, no relation to the author.

(33) The Mind's New Science, 12.

(34) Logical positivism, in a vindication of Newton's laws of motion or of Hegelianism perhaps, inspired its own opposite in the form of phenomenology. Martin Heidegger, whose work would become more influential through the second third of the 20th century, did not proclaim the importance of consciousness in a vacuum. On the contrary, he was speaking in direct response to the ongoing insistence on the sole primacy of formal statements of existence.

(35) This is a postdoctoral thesis, which remains a common prerequisite for tenure privileges in universities on the European Continent.

(36) Blackburn, Oxford Dictionary of Philosophy.

(37) Stoll 1961, 167.

(38) On the contrary: "Godel's procedure and his results opened the way to a fully rigorous treatment of the notion of a computable function and to our modern understanding of the power and limits of computation, and the possibility or otherwise of programs that test for consistency and completeness." Blackburn, Oxford Dictionary of Philosophy.

(39) J. Symbolic Logic, vi, 1936, 103-105.13. "Finite Combinatory processes, Formulation 1". J. Symbolic Logic, vi, 1936, 103-105, Post.

(40) "On Computable Numbers", Proc. London Math Soc, ser.2, v 42 1936, 230-265.FN 16: "On Computable numbers with an application to the Entscheidungsproblem", Proceedings of the London Mathematical Society 42 (1937).

Noted science writer Charles Petzold, has written The Annotated Turing: A guided tour through Alan Turing's historic paper on

Computability and the Turing Machine, an extraordinary edited presentation of this paper.

(41) Ibid.

(42) Aspray 1990, 117.

(43) 1990 117.

(44) Cite new biog of Wiener.

(45) 1943, 18.

(46) Ibid, 20.

(47) Philosophy of Science, January 1943, 23.

(48) Blake and Uttley 1959, 92; Heims 1991.

(49) Ref to computer history during the 1940s.

(50) Shannon spent his career as an MIT professor and sometime Bell Labs visitor. Vannevar Bush was not related to the United States presidents of the same surname. See Z. Pascal's thorough biography, and Ceruzzi and Aspray histories of the computer.

(51) Analog computer: a computer that measures continuously changing conditions, such as temperature and pressure, and converts them into quantities (Webster's NewWorld Dictionary of Computer Terms).

This computer, like many analog machines, worked quite well. But from our current perspective it was too dependent on the physical stuff of its materials, and this limited the precision of its calculations.

(52) Evans 1981, 60.

(53) Shannon and Weaver 1949, 8.

(54) Shannon and Weaver 1949, 20.

(55) Lucky 1989, 41.

(56) Lucky 1989, 58.

(57) In addition to being the catalyst for the digital computer, this project was critical for the invention of Cybernetics' concepts, as discussed in the last chapter.

(58) Zachary, Endless Frontier, 1.

(59) Dickson 1971, p13-14.

(60) Josef Stalin putatively endorsed the concept of 'Socialism in one country', in contrast to the internationalism of Lenin. Unfortunately, the definition of 'country' was slippery, including dozens of satellite states and the incorporation of numerous separate provinces into the Soviet Union proper.

(61) Rhodes, The Making of the Atomic Bomb, 19.

(62) Bush, Science, the Endless Frontier, 22.

(63) Bush 1945; Pennick et al., 1973.

(64) The National Cancer Institute had been established in 1937, and the National Institutes of Health had been bolstered during the war.

(65) Atomic energy was the only branch of science that was unilaterally exempt from Congress' suspicion of 'undemocratic' expenditures. The Atomic Energy Commission was created in 1945, and this field's expressly national laboratories- including Sandia, Livermore, and Los Alamos- were spared further scrutiny. Penick et al. 1972, p17; Foerstel, Secret Science, 1993. See also Fred Inglis, The Cruel Peace: Everyday Life and the Cold War, 1991.

- (66) Penick et al. 1972, 22.
- (67) Penick et al. 1972, 23.
- (68) Penick et al. 1972, 17.
- (69) Norberg and O'Neill 1996, 4.
- (70) Dickson Think Tanks, 1971, 23.
- (71) Dickson Think Tanks, 1971, 23.
- (72) Leslie, *The Cold War and American Science*, 8.
- (73) Hapgood 1993; Leslie, *The Cold War and American Science*.
- (74) Penick, 1972.
- (75) Norberg and O'Neill, 1992, vii; Aker 2008.
- (76) Leslie *The Cold War and American Science*, p25.
- (77) Hapgood, 1993.
- (78) Licklider, CBI Interview 150, 1988; Waldrop 2001.
- (79) For instance, J.C.R. Licklider recruited William McGill, then-future president of Columbia University, into his post-graduate projects.
- (80) Winston CBI interview, 1990.
- (81) Shasha and Lazere, *Out of their minds: The Lives and discoveries of 15 great computer scientists*.
- (82) Hilts 1983, 217-218.
- (83) Bernstein, 38; Blake and Uttley 1959, 4. Minsky has also maintained a singularly mixed collection of interests. Bernstein

recounts his abortive effort, around 1970, to market a cheap, popular general-usage computer (p22). He has also dabbled in optics, and invented a device which sensed distinctions in vocal tones (Bernstein 1983, p72), consulted to the movie 2001, and written science fiction.

(84) Stanford's Medical School was located in San Francisco at that time. This and the remainder of biographical details in this paragraph are from Newell's CBI Oral History.

(85) Crowther-Heyck, Hunter. Herbert A. Simon: The Bounds of reason in Modern America, 17.

(86) MOML p83.

(87) Bernstein, 27.

(88) MOML, p125-132.

(89) Newell CBI OH.

(90) Crowther-Heyck.

(91) Simon, MOBR article regarding the model of the Servomechanism.

(92) Edwards, The Closed World.

(93) Heims 1991; M. Waldrop's 2003 biography.

(94) Licklider CBI OH interview.

(95) Computing evoked keen interest among the science policy uber-elite. Dr. Mina Rees, president of the graduate school at CUNY and later deputy chief scientist at the ONR, was one of the early visitors to Goldstine and co., and an early supporter in science policy circles, Goldstine, The Computer from Pascal to Von Neumann, p212.

(96) Goldstine, The Computer from Pascal to Von Neumann, p275. The group also included S.S. Wilks of Princeton, Walter H. Pitts, who was at the time at the Kellex Corporation instead of U. Chicago, E.H. Vestine of the Carnegie Institute of Technology, E.W. Deming from the United States Census; Warren McCulloch of the U. of Illinois Medical School, Lorento de No of the Rockefeller Institute, L.E. Cunningham, Director of the BRL, and Goldstine himself.

(97) Heims, John von Neumann and Norbert Wiener: From Mathematics to the Technologies of Life and Death. Also see McCorduck, Machines who think, 78-79.

(98) Uttley, as a Telecommunications Research Establishment engineer, designed the 1948 TRE digital computer. Goldstine, The Computer from Pascal to Von Neumann, 218.

(99) Waldrop, The Dream Machine: J.C.R. Licklider and the revolution that Made Computing Personal. Also, JCR CBI oral history.

(100) Jeffress, ed. Cerebral Mechanisms in Behavior: The Hixon Symposium.

(101) Heims.

(102) The Mind's New Science discusses this topic at length.

(103) Grosch, Computer: bit slices from a life.

(104) In 1943, Pitts and McCulloch, had written:
" Anything that can be exhaustively and unambiguously described [in logic]... is...realizable by a suitable finite neural network."

McCulloch and Pitts, " A Logical Calculus of the Ideas Immanent in Nervous Activity". Bulletin of Mathematical Biophysics. 1943.

Their words proposed that computers could be built with biological rather than only physical inorganic components. The

discussion was highly theoretical rather than practical, but its erudition was such that that was hardly the point. The 'neurons' could be designed to emulate Boolean logic gates (i.e., logical functions), and hence to compute anything that could be stated in a logical formalism. Their work does not offer immediately evident opportunities for digital computing, which was made up of metal and bulbs and magnets. Yet it did offer the spectre of an artificial neural system for sensing things, seeing things, perceiving distributed stimuli, etc. As such, this work and the continued enthusiasm of McCulloch in particular, inspired the era's artifacts and work in automata.

The model of information transmission suggested in the 1943 paper was, in retrospect, found to be incorrect. It was later discovered that neurons were quite distinctly differentiated according to function, rather than being uniform as assumed. Moreover, it was discovered that they did not act according to a simple auction model, in which a given level of stimulation would release an excitatory or inhibitory firing. This had been the contemporary wisdom, proposed by Sherrington, circa 1940. But the suggestion that neurons, real or artificial, 'thought' in a way that was readily isomorphic to logic made it seem possible. The effort inspired original research in artificial vision, pattern recognition, haptics, and other forms of sensory emulation.

Pitts and McCulloch are discussed in Heims 1991; McCorduck 1979 p78-79, and Papert in McCulloch 1965.

(105) Goldstine 1972, p274. Also see Aspray, John Von Neumann and the Origins of Modern Computing.

(106) Martin Campbell-Kelly and William Aspray's Computer: A History of the Information Machine discusses the EDVAC report's messy publication and public presentation. 2004.

(107) Goldstine 1972, p276.

(108) Evelyn Fox Keller's article, "Marrying the Premodern to the Postmodern: Computers and organisms after World War II, in

Franchi and Guzeldere's edited volume 2005, examines this provenance in detail.

(109) Von Neumann." The General and Logical theory of automata," In Pylyshyn, Zenon. ed. Perspectives on the Computer Revolution. Englewood Cliffs, N.J.: Prentice-Hall. 1970: 87-113. p98.

(110) Von Neumann 1948 paper, 94.

(111) Von Neumann 1948, 101.

(112) " Had he been able to bring the power of analysis to bear on formal logics and automata theory, Von Neumann's results would certainly have been of the greatest interest. In particular, Burks conjectured that he thought of differential equations in respect to his excitation- threshold- fatigue model. The problem then, is largely connected w the behavior of neurons when stimulated. This brings us very close to the brilliant work of Alan L. Hodgkin and Andrew F. Huxley, for which they received the Nobel Prize for medicine in 1963. They described the behavior of nerve fibers by means of a non- linear partial differential equation." Goldstine 1972, p279.

(113) This was not a surprising fate for anyone who had witnessed, at relatively close range, the Alamogordo blast of July 16, 1945. Incredibly, the main safety measure at this event was averting one's eyes. Richard Rhodes, The Making of the Atomic Bomb.

(114) Von Neumann 1958, ix.

(115) Turing was homosexual at a time when homosexuals were actively persecuted in Britain and elsewhere. The torture disguised as hormonal 'treatments' he was forced to undergo led to painful and emasculating physical side effects. See A. Hodges' biography, Turing: the Enigma of Intelligence.

As B. Jack Copeland, editor of *The Essential Turing: The ideas that gave birth to the Computer Age*, points out, Turing was subject to 'the shabbiest of treatment from the country he had helped save', p12.

Also see the more recent biography by T. Gottfried, *Alan Turing: The Architect of the Computer Age*, 1996; and Dewdney, *A Turing Omnibus*.

(116) Turing article in Bowden ed.; McCorduck.

(117) "In the correspondence between Turing and the biologist J.Z. Young, it is clear that Turing by then knew of Warren McCulloch's work in the U.S. and was likewise convinced that a mathematical approach was more fruitful than an anatomical one to the problem of brain function." McCorduck, *Machines Who Think*, 59.

(118) When Baumgartner and Sabine interviewed prominent figures in AI and cognitive science for *Speaking Minds: Interviews with Twenty Eminent Cognitive Scientists* (1995), a number of the interviewees answered that they were not very interested in the Turing Test, and rarely thought about it.

(119) Turing 1950, 19.

(120) Turing 1950, 19.

(121) Turing 1950, 16.

(122) Bernstein, *Science Observed*.

(123) Apparently the feeling was mutual: B.F. Skinner's autobiography does not mention any of the AI people either.

(124) Bernstein *Science observed: essays out of my mind*, 30.

(125) Bernstein Science observed: essays out of my mind, 30. Also see Marvin Minsky and Seymour Papert, Perceptrons: an introduction to computational geometry, px.

(126) The Snarc was mysteriously spirited away to Dartmouth sometime during the 1950s, and disappeared. Bernstein; Science observed, 36-37.

(127) Minsky 1956: 117-128.

“In the present paper it is shown that a certain category of sets of elements are universal in the sense that one can assemble such elements into machines with which one can realize functions which are arbitrary to within certain reasonable restrictions.” Minsky 1956, 117.

(128) “ The hypothesis the con junctions and dis junctions are initially available is motivated by the prevailing opinion in neurophysiology that such elements are almost certainly represented among and in fact are probably characteristic of the cells of the CNS. On the other hand the nature and distribution of the non-monotonic properties of the nervous system are not nearly so well understood, in particular the various forms of inhibition. Thus the particular form of the theorem may be of some value in analyzing those neural phenomena in which there appears to be an inhibitory quality but in which no specific inhibitory connection or mechanism has been isolated.” (Minsky 1956, p127).

(129) Hilts 1983, p215-216.

(130) C. Shannon and J. McCarthy, "Automata Studies", Annals of Mathematical Studies N34, Princeton, N.J.: Princeton University Press. 1956)

(131) McCorduck, Machines who Think, 102.

(132) McCarthy, CBI interview 156.

(133) Newell CBI OH.

(134) McCarthy 1957, 177.

(135) John McCarthy, "The Inversion of Functions Defined by Turing Machines," In Shannon, C. and J. McCarthy, eds. "Automata Studies", Annals of Mathematical Studies N34, Princeton, N.J.: Princeton University Press. 1956: 177-182.

(136) Von Neumann and Morgenstern, game theory work-publisher ref.

(137) Abella, *Soldiers of Reason*, p9.

(138) Openly criticizing Von Neumann's cynical statement that human game-players are rational and perfectly informed nuclear war conceived as a game, Norbert Wiener maintained his honesty at the expense of being popular. See Heims 1980, p313-315. Conway and Siegelman discuss this at length as well; p253.

(139) Computers that controlled weapons systems, and envisaged wars using weapons as a game, have haunted the dreams of sci-fi movie screenwriters. As the WOPR, or War Operation Plan Response computer's says in its final assessment after testing out a series of war scenarios in War Games (1983), "The only winning move is not to play." See Kaplan, *The Wizards of Armageddon*.

(140) NSS 1958, p42-44, has an excellent explanation of the concept.

(141) Shannon: This article was apparently not connected to Caissac, Shannon's chess-playing machine (Lucky 1989).

(142) Bowden, *Faster than Thought*, 1953; NSS 1958, p44.

(143) NSS 1958, p48.

(144) NSS 1958, p45.

(145) McCorduck, *Machines Who Think*, p159.

(146) McCorduck, *Machines Who Think*, 159.

(147) NSS 1958, 48.

(148) McCorduck, *Machines Who Think*, 149.

(149) Computer impresario Louis Ridenour had a hand in various pursuits in early digital computing. He encouraged Warren Weaver in his publications with Shannon, and was engaged in various burgeoning electronics business. He later became a vice president of Lockheed.

(150) "It happened there was to be a world checker champion meeting in neighboring town of Kankakee, somebody got the idea- I'm not sure it was mine, but I got blamed w it at least- that it would be nice to build a small computer that could play checkers. We thought checkers was probably a trivial game. Claude Shannon had talked about programming a computer to play chess... we decided to pick a simpler game... then at the end of the tournament we'd challenge the world champion and beat him, you see, and that would get us a lot of attention. We were still very naive..." McCorduck, *Machines Who Think*, 149.

(151) McCorduck, *Machines Who Think*, p150.

(152) A. Samuel. "Some studies in machine learning using the game of checkers" IBM Journal of Research and Development July, 1959. Reprinted in Computers and Thought.

(153) B. Jack Copeland, ed., *The Essential Turing*, 355-258.

(154) Samuel 1959, p72.

(155) Grosch, *Computer: Bit Slices from a Life*.

(156) McCorduck, *Machines Who Think*, p152.

(157) Smithsonian Rand videohistory, Shaw and NP35 RU9536 session 6.

(158) Allen Newell, "The Chess Machine: an Example of Dealing with a Complex Task by Adaptation". Proceedings of the 1955 Western Joint Computer Conference. Institute of Radio Engineers, NY: 101-107.

See also:

Newell, A., J.C. Shaw, and H.A. Simon." Chess-Playing Programs and the Problem of Complexity". IBM Journal of Research and Development 2, 4. October, 1958: 320-335. Reprinted in Computers and Thought.

(159) Newell 1955, p108.

(160) Chapters 10 and 11 of Crowther-Heyck's Herbert Simon discuss the appearance and arduous cultivation of the analogy between programs and minds.

(161) Following the war, personal acrimony divided Norbert Wiener and his longtime colleague Warren McCulloch. Conway and Siegelman's biography, rich in psychological portraiture, explains that Wiener turned his back on McCulloch, Pitts, and computing in general- to the loss of all. Wiener wrote increasingly angst-ridden essays on the dangers of the new nuclear sciences and computers. It is indeed curious that one of the single brightest scientists of the 20th century read technological progress as an unlimited disaster, and became a veritable Luddite:

" The automatic factory and the assembly line... gives the human race a new and most effective collection of mechanical slaves to perform its labor".

Cybernetics, quoted in Conway and Siegelman, Dark hero of the Information Age: In search of Norbert Wiener, the father of Cybernetics.

(162) Berkeley, Giant brains; or, Machines that Think, 7, 182.

Edmund Callis Berkeley was one of the earliest believers in digital computing. As an actuary working at Prudential Insurance of New

York City in the late 1940s, he tried unsuccessfully to persuade his employer to buy a Univac. In 1947, he called the meeting at Columbia University at which the ACM was formed. The next year, he established Computers and People, an early popular magazine on computing. Berkeley continued to write popular works on computing, as well as specialized ones, such as a LISP language manual which he and Daniel Bobrow (later a major AI figure) wrote in 1964. As we saw in Chapter Four, Berkeley also created robotic automata.

(163) Berkeley 1949, p7.

(164) Bowden 1953, p317.

(165) Bowden 1953, p319.

(166) Bowden 1953, p320.

(167) Turing, A.M." Computing Machinery and Intelligence". (1950). In Feigenbaum, Edward A. and Julian Feldman, eds. Computers and Thought. New York: McGraw-Hill. 1963: 11-39.

(168) Ibid. Also see Turing, A.M., "Intelligent Machinery". (1947). Machine Intelligence 5. B Meltzer and D Michie, eds. Edinburgh: Edinburgh University Press. 1969.

(169) A dozen years later, a literature review by Paul Armer, Director of Computer Sciences at RAND, would outline similar arguments in his literature review, "Attitudes toward intelligent machines"; 1963 compendium Computers and Thought.

(170) Turing 1950, p20.

(171) Turing 1950, p27.

(172) Turing 1950, p24.

(173) Turing 1950, p26-27.

(174) Martin, C. Dianne. "The Myth of the Awesome Thinking Machine", Communications of the ACM 36, 4 (April, 1993): 120-133.

(175) Bowden 1953, vii.

(176) McCorduck, *Machines Who Think*, p44; refers to Boring, E. "Mind and Mechanism", American J of Psychology 2, April 1946.

(177) McCorduck, *Machines Who Think*, 83.

(178) The RAND Corporation: the First Fifteen Years. Santa Monica, CA. 1963.

(179) The RAND Corporation: the First Fifteen Years. Santa Monica, CA. 1963.

(180) See Ch. V in Dickson, Think Tanks, and Rand's own webpage.

"At the end of 1945, "General H.H. "Hap" Arnold, Commanding General of the Army Air Forces... suggested and effected a contract between the Army Air Forces and the Douglas Aircraft Company." Ten million dollars was originally allocated for the project. *Ibid.*, 6.

(181) Dickson, *Think Tanks*, and The RAND Corporation.

(182) See Alex Abella's wonderful *Soldiers of Reason*, as well as Kaplan's *Wizards of Armageddon* and Abbate's *Inventing the Internet*. Several years following our time frame, Rand researcher Paul Baran would conceive of the Internet as an outline of proposed post-nuclear holocaust electronic communication.

(183) Abella, 94.

(184) The RAND Corporation: the First Fifteen Years, 1.

(185) Heims, John von Neumann and Norbert Wiener: *From Mathematics to the Technologies of Life and Death*, 314.

More generally, see Baum, *The System Builders*.

(186) Mission-directed research is concerned with developing particular pieces of technology, such as missiles or number-crunching computers. This differs from basic research, for instance the understanding of the working of the body or of physical forces, which is by definition not immediately manifest in a technology.

(187) From *Science*, in a satirical interview with a “Dr. Grant Swinger,” identified as the director of ‘Breakthrough Institute and Chairman of the Board of the Center for the Absorption of Federal Funds’. Dickson, *Think Tanks*, 39.

(188) During the mid-Twentieth century, work outside the normal ‘business hours’ was so odd that it is even discussed as an innovative practice in *The RAND Corporation: the First Fifteen Years*.

(189) Smithsonian Interview on Rand Corporation. RU536 Rand, P65, session 5, Willis Ware.

(190) SAGE was an enormous endeavor; see Kent Redmond and Thomas Smith *From Whirlwind to MITRE: the R&D Story of the SAGE Air Defense Computer*. MIT Press, 2000; and Philip Mirowski’s *Machine Dreams*.

(191) Edwards, *The Closed World*, 121-122.

(192) McCorduck, *Machines Who Think*, 117.

(193) Crowther-Heyck, Herbert A. Simon, 203.

(194) MOML, 200.

(195) Crowther-Heyck, Herbert A. Simon, 205.

(196) Newell interview, CBI Oral History Interviews.

(197) Smithsonian Interview on Rand Corporation. RU536 Rand, P65, session 5, Willis Ware.

(198) Ibid.

(199) McCorduck, Machines who Think, 127.

(200) Ibid. MOML, 170.

(201) McCorduck, Machines Who Think, 125.

(202) MOML, 201.

(203) McCorduck, Machines Who Think, 132.

(204) MOML, 201.

(205) Smithsonian RAND Videohistory, RU9536 Session 6 RU9536, 31-33 quote from Shaw.

(206) MOML, 201.

(207) Shaw is an excellent example of the superior opportunities offered almost exclusively to white men from most socioeconomic backgrounds as a result of the economic and military boom of the PostWar period. Shaw's family ran a paint store in a small Southern Californian town, and he learned navigation during the war. Shaw worked as an actuary for an insurance company until he moved up to Rand, in 1950 (Smithsonian Institution, Clifford Shaw papers). As we have seen in Chapter 6, he was one of the inner circle of AI in its infancy.

(208) Newell CBI.

(209) Smithsonian RAND Videohistory, RU9536 Session 5 and 6.

(210) MOML, 170.

(211) Ibid.

(212) Newell CBI, MOML p168. Also see Baum, The System Builders and The Rand Corporation.

(213) Having elected American citizenship, he spent a year in the Navy "I got my commission as an ensign five days after VJ day. I spent the next year mostly in delight around the Pacific. Then I returned to graduate school at MIT under the GI bill." Email communication with OGS, February 11 2003.

(214) Blake and Uttley 1959, p512; personal communication with the author.

(215) McCorduck, Machines Who Think.

(216) Bernstein 1983, p72-73; personal communication with the author.

(217) Heims 1991, p44; Conway and Siegelman.

(218) Semantic Information Processing.

(219) Newell, Allen. Interview by Arthur L. Norberg, audio tape(s) and transcript, Pittsburgh, Pa., 10-12 June 1991, Charles Babbage Institute Interview. OH 227.

(220) Selfridge was also elected an AAAI Fellow in the mid-1980s, and gave talks at AI conferences over the years; Newell CBI OH.

(221) McCorduck, p135.

(222) Newell, Allen. Interview by Arthur L. Norberg, audio tape(s) and transcript, Pittsburgh, Pa., 10-12 June 1991, Charles Babbage Institute Interview. OH 227.

(223) Polya, How to Solve It, 114. As we saw, Polya was one of Newell's professors at Stanford. Polya Hall, the original site of Stanford's Computation Center, was named after Professor Polya.

See also Peter Duren, ed., A Century of Mathematics in America, p273.

(224) Crowther Heyck, 224.

(225) McCorduck, Machines Who Think, 137.

(226) Principia Ch 2, 2.15 is the exact proof.

(227) MOML, 206.

(228) Simon 1991. Also Smithsonian Rand VideoHistory RU 9536 Session 6, P34 Shaw.

Logic Theorist: 1957, reprinted 1963; Newell and Simon 1972, Simon 1991 p207-29; and Handbook of AI, vi, p109-111.

(229) McCorduck, Machines Who Think, 142.

(230) Goldstine 1972, 309.

(231) MOML, 205.

(232) McCorduck 1979, 132 and MOML and Feigenbaum paper.

(233) McCorduck 1979, 139.

(234) Simon MOML.

(235) MOML, 206.

(236) MOML, 206-207.

(237) MOML; Simon indicates that this story comes from Feigenbaum, not Simon himself.

(238) J. McCarthy; M.L. Minsky, N. Rochester, and C.E. Shannon." A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence".

(239) Bernstein, Science observed: essays out of my mind, 39-40.

(240) Denicoff, " AI Development and the Office of Naval Research". In Bartee, T., ed. Expert Systems and Artificial Intelligence: Applications and Management, 271-289.

(241) The final list is immediately above Notes.
Drawn from the JOHN MCCARTHY WEBSITE:

The final list of people who attended, visited, or were invited to the conference is:

Marvin Adelson (Hughes Aircraft Company, Los Angeles);
Peter Elias (R. L. E., MIT);
W. Duda (IBM Research Laboratory, Poughkeepsie, NY);
Paul Davies (Los Angeles, CA);
B.G. Farley (Arlington, MA);
E.H. Galanter (University of Pennsylvania);
Herbert Gelernter (IBM Research, Poughkeepsie, NY);
Harvey A. Glashow (Ann Arbor, MI);
Herbert Goertzal (New York, New York);
Leon D. Harmon (Bell Telephone Laboratories, Murray Hill, NJ);
John Holland (E. R. I. University of Michigan at Ann Arbor);
Anatol Holt (Philadelphia, PA);
William H. Kautz (Stanford Research Institute);
R.D. Luce (New York, NY);
Z.A. Melzak (Mathematics Department, University of Michigan at Ann Arbor);
Trenchard More (Department of Electrical Engineering, MIT);
Abraham Robinson (Department of Mathematics, the University of Toronto);
Hartley Rogers, Jr. (Department of Mathematics, MIT);
Jerome Rothstein (Red Bank, NJ);
David Sayre (IBM Corporation, New York, NY);
J.J. Schorr-Kon (Lincoln Laboratory, MIT);
L. Shapley (Rand Corporation);
M.P. Schutzenberger (R.L.E., M.I.T.);
Raymond J. Solomonoff, Technical Research Group (New York, NY);

J.E. Steele, Capt. USAF (Wright-Patterson AFB, Ohio);
Frederick Webster (Cambridge, MA);
E.F. Moore (Bell Telephone Laboratory, Murray Hill, NJ);
and John Kemeny (Dartmouth College).

Oddly, this list lacks the name of Arthur Samuel, who did attend the conference (McCorduck).

(243) Newell and Simon, Human Problem Solving, 884.

(244) Newell and Simon, Human Problem Solving, 884. Bernstein, Science Observed.

(245) McCorduck, Machines who think, 106.

(246) McCorduck Machines who think. Simon did not press this point in MOML.

(247) Newell CBI OH interview 1991.

(248) Herbert A. Simon personal interview with the author, April 10, 1997.

(249) McCorduck, Machines who think, 106.

(250) McCorduck, Machines who think, 101.

(251) Simon, Models of My Life, 210.

(252) Proceedings, 108.

(253) Gardner's The Mind's New Science discusses the publication of this paper in 1956 as a seminal contribution to the Cognitive Revolution; 89.

(254) Gardner explains Chomsky's earliest work in Chapter Seven of The Mind's New Science.

(255) This individual was not Arturo Rosenblueth, the distinguished Cyberneticist who coauthored "Behavior, Purpose and Teleology" and was a longtime participant in the Macy Conferences. Walter Rosenblith was a physiologist who wrote papers on the nature of electrical impulses in the CNS. See Feigenbaum and Feldman eds., Computers and Thought 1995, 512.

(256) Simon, Models of My Life, 210-211.

(257) H.A.S. letter to RES, March 10, 1997.

(258) Simon MOML.

(259) Newell CBI interview,

(260) Newell and Simon's work with psychologists appeared in the volume Representation and Meaning.

(261) MOML, 218.

(262) In the 1990s, EPAM was still being used as a research tool. Personal communication, Pat Langley, January 2001.

(263) Herbert A. Simon, Donald W. Smithburg, and Victor A. Thompson. First edition, New York, Knopf, 1950. Administrative Behavior by Herbert A. Simon, New York, Macmillan Co. 1947.

(264) H.A. Simon, Models of man: social and rational; mathematical essays on rational human behavior in a social setting. New York, Wiley, 1957.

(265) McCorduck 1979, 147.

(266) "Rational choice and the structure of the environment", 1956.

(267) "The behavioral theory of the firm", 1963 paper.

(268) Herbert A. Simon, "A Behavioral Model of Rational Choice", Quarterly Journal of Economics 69: 99-118.

(269) Crowther-Heyck, Herbert Simon, 127-131.

(270) It is a testimony to HAS' multifaceted contributions to economics and other fields that treatments of his other academic contributions fill festschriften and entire books on single fields. RES cannot discuss this further, but HAS continues to inspire excellent scholarship. Those which most notably come to mind are Mie Augier's articles, the Klahr and Kotovsky festschriften, the Augier and March festschriften, Crowther-Heycke's biography of Simon, and Philip Mirowski's masterpiece, Machine Dreams.

(271) Bruner, Goodnow, and Austin 1956, pviii.

(272) Bruner, Goodnow, and Austin 1956, vii.

(273) Bruner, Goodnow, and Austin 1956, ix.

(274) Jean Piaget's original field of study was zoology. Surely this is a symptom of the paucity of studies of cogitation in the early part of the 20th century.

(275) MOML, p190.

(276) Crowther-Heycke, 243.

(277) Drosophila is 'any of a genus of small two-winged flies used in genetic research'. That is, drosophila are the fruit flies one reads of in reports on biology experiments.

(278) Clifford Shaw's participation in this project waned. Most publications on GPS, unlike those on LT, reflect only the two authors and later co-author participants. This is not entirely the case- for instance, "The process of Creative Thinking" in Models of Thought (1962), is an NSS work. But Newell and Simon's 1972 oeuvre Human Problem Solving is dedicated to Shaw.

(279) NSS 1962. The first paper was, 'The Processes of Creative Thinking', mentioned above. MOML, p221.

(280) The most important papers covering the GPS program[s] are the 1961 article explaining the reasoning behind the GPS (the 1961 article was reprinted in Computers and Thought), their 1959 report to a Paris UNESCO conference regarding GPS' performance on solving trigonometric problems, and Ernst and Newell's final report in 1969. Models of my Life, Machines Who Think, Crevier's history of AI, and Volume III of the Handbook, all also provide technical analysis and 'in-depth' analysis by the program's creators in context of the era. Newell also offers an interesting sub-note on GPS on pages 227-228 of Unified Theories of Cognition (1991).

(281) A protocol is a log of the spoken-out problem-solving process as voiced by a test subject. In this case, test subjects were CIT undergraduates. The term and the practice were borrowed from deGroot's study of chess-playing, as mentioned earlier, as well as from work being done at that time by Naval Research Laboratory psychologists O.K. Moore and Scarvia B. Anderson. McCorduck Machines Who Think, 212.

(282) Charniak and McDermott AI Textbook 1985, 301.

(283) In the contemporary classroom, 'Cannibals and Missionaries' has been renamed 'Wolves and Goats' to avoid offending cannibals.

(284) Charniak and McDermott 1985, 302.

(285) Minsky, 1979.

(286) Handbook of AI I, 56-57, 59-60.

(287) Handbook of AI I, 56-64.

(288) MOML, 166.

(289) Newell, A., J.C. Shaw, and H.A. Simon." Chess-Playing Programs and the Problem of Complexity". IBM Journal of Research and Development 2, 4. October, 1958: 320-335. Reprinted in Computers and Thought.

(290) NSS 1958, 51.

(291) Ibid., 63.

(292) NSS 1958, 8-9.

(293) Hapgood 1993.

(294) Leslie, *The Cold War and American Science: The Military-Industrial-Academic Complex at MIT and Stanford*, 15.

(295) Regarding MIT's transformation, see also Henry Etkowitz, MIT and the rise of entrepreneurial Science, 2002.

(296) A. Aker, Calculating a natural world, addresses this secular trend, as impressive in scale as the war effort itself. Waldrop's *Dream Machine* (2001) is an extremely engaging biography of J.C.R. Licklider and also of his times, in which MIT played an immense role.

(297) The hard-science tone of MIT extended itself to the social sciences there as well. The linguistics department, led by Noam Chomsky for decades, has helped to foster the idea of innate mental and linguistic facilities.

(298) Leslie, *The Cold War and American Science: The Military-Industrial-Academic Complex at MIT and Stanford*, 20.

(299) Dickson, *Think Tanks*, 153.

(300) Leslie, *The Cold War and American Science: The Military-Industrial-Academic Complex at MIT and Stanford*, 15.

(301) Pfeiffer 1962, 1.

(302) Fernando J. Corbato, CBI Interview 162, 1989.

(303) Kahn, CBI OH 192, 1990.

(304) Minsky, CBI OH 179, 1989.

(305)Blake and Uttley 1959, p4.

(306) Bernstein 1983, p72.

(307) LIST OF PROGRAMS CARRIED OUT at MIT- before 1961:
- N. Rochester, " Symbol Manipulating Language." 1958;
- L. Hodes," Some results from a pattern recognition program using LISP." n.d., prob circa 1958;
- J. Slagle, SAINT (symbolic automatic integrator); 1961.

(308) of both the Logic Theorist and Minsky's ideas (e.g., Gelernter and Rochester 1958).

(309) John McCarthy, Aspray interview CBI-OH 156),

(310) McCarthy 1983.

(311) Or so the accounts go. However, the Uttley Conference biographical data on McCarthy says that he became an assistant professor of communication sciences (Blake and Uttley 1959, p76).

See re Minsky: " Machines with Experience", Time Magazine, volume 72, n23 (December 8, 1958): p73. Also see The New Yorker (New York, N.Y.: 1925). THE NEW YORKER ([New York : F-R Pub. Corp.); spec: Dec 6 1958; 44-45.

(312) Kemeny later became the president of Dartmouth as well. Bernstein 1983, p75.

(313) Sources concerning the early development of LISP include the following:

John McCarthy, " Recursive Functions of Symbolic Expressions" In Horowitz, ed. Programming Languages: A Grand Tour. Rockville, MD: Computer Science Press. 1987: 174-202.

McCarthy, John; Paul Abrahams, Daniel Edwards, Timothy Hart, and Michael I. Levin. LISP 1.5 programmer's manual. The Computation Center and Research Laboratory of Electronics, Massachusetts Institute of Technology. Second edition. Cambridge, MA: The MIT Press. 1965.

Herbert Stoyan, " Early LISP History (1956-1959)". University of Erlangen-Nurnberg Lehrstuhl fur Kunstliche Intelligenz Am Weichselgarten 7, D-91058 Erlangen, Germany.

www.informatik.uni-erlangen.de

McCarthy himself has written several articles on the topic of LISP history: for instance, McCarthy, John." History of LISP". SIGPlan Notices 1978. 13: 217-223; and an article in R.L. Wexelblatt, ed. History of Programming Languages. New York: Academic Press. 1981.

The more recent development of LISP is well-documented, although it is usually presented in the form of textbooks (e.g., The Little Lisper) rather than addressing the topic of the language as such or in historical context. One very thoughtful analytical article is:

P. Wegner," Programming Languages: the First Twenty-Five Years". In Horowitz, ed. Programming Languages: A Grand Tour. Third Edition. Rockville, MD: Computer Science Press. 1987: 4-23. More generally, Naomi Baron's book, Computer Languages: A Guide for the Perplexed (1986) is also helpful.

- Newell, Allen, ed. (The Rand Corporation). Information Processing Language V Manual. Englewood Cliffs, N.J.: Prentice Hall. 1961. **owncopy - Newell 1961;- Newell and other articles in Yovits 1962; - see McCracken 1957;- from Aspray- Campbell Kelly integrated.

(314) Wegner 1976, p9.

(315) lambda notation: a term such as 'sine 60' or 'the father of Hegel' refers to a number or a person. But it includes the term 'sine...' or 'father of...'. 'Sine x' or father of y' stand for a function referring respectively to a number or a person for particular

values of x and y . If we wish to refer to the function itself, the notation $(\lambda x)(\text{father of } x)$ or $(\lambda y)(\text{father of } y)$ is adopted. A logical calculus with rules involving such terms is called the lambda calculus." Blackburn, The Oxford Dictionary of Philosophy.

(316) Charniak and McDermott 1985; Wilensky 1984.

(317) Stoyan 1998; Wegner 1976).

(318) John McCarthy, "Programs with Common Sense". In Blake, D.V. and A.M. Uttley, eds. Proceedings of the Symposium on the Mechanisation of Thought Processes. National Physical Laboratories, Teddington, England. London: H.M. Stationary Office. 1959: 75-84. Reprinted in Minsky ed., Semantic Information Processing. Cambridge, MA: The MIT Press. 1968. The Advice Taker paper has been reprinted at least twice: Minsky ed., Semantic Information Processing. Cambridge, MA: The MIT Press. 1968. McCarthy, J." Programs with Common Sense." In Luger, George, ed. Computation and Intelligence: Collected Readings. Cambridge, MA: The MIT Press. 1995.

(319) McCarthy 1959, p80-82.

(320) McCarthy 1959, p80-82.

(321) "... I do not think there could possibly exist a programme which would, given any problem, divide all facts in the universe into those which are and those which are not relevant for that problem. Developing such a programme seems to me to be by 10 to the 10 orders of magnitude more difficult than say the Newell Simon problem of developing a heuristic for deduction in the propositional calculus. This cavalier way of jumping over orders of magnitude only tends to becloud the issue and throw doubt on ways of thinking for which I have a great deal of respect. By developing a powerful programming language you may have paved the way for the first step in solving problems of the kind treated in your example, but the claim of being well on the way

towards their solution is a gross exaggeration. This was the major point of my objections.” (McCarthy 1959, p88).

(322) Microsoft Press Computer Dictionary. Second Edition. 1996.

(323) Peterson, Institute Historian T. F. Nightwork: A History of Hacks and Pranks at MIT. The MIT Press. 2003.

(324) Oliver G. Selfridge, email to RES, 11 Feb 2003 12:03:11 -0500).

(325) Levy, Hackers.

(326) Hapgood 1997 and Levy 1984 describe this history very well.

(327) Levy, Hackers.

(328) Klahr and Kotovsky 1992, p166.

(329) Larry Roberts CBI OH interview.

(330) Levy, Hackers, 23.

(331) Hapgood 1997.

(332) Levy, Hackers.

(333) Fano 1964.

(334) Williams 1979, 399.

(335) The movement of information between magnetic core and drum was called “automatic” swapping, meaning that the machine had virtual memory. Prior to this invention, virtual memory had meant that the programmer had had to look up the relevant information on the drum.

(336) Michael Williams' A History of Computing Technology observes: " Had this been incorporated into the machine, we would likely have seen the mass produced time-shared computer being commercially available a few years earlier than it actually was." 402.

The movement of information between magnetic core and drum was called "automatic" swapping, meaning that the machine had virtual memory. Prior to this invention, virtual memory had meant that the programmer had had to look up the relevant information on the drum.

(337) Campbell-Kelly and Aspray indicate that Strachey patented this idea as well, 1996, p209.

(338) McCarthy 1959, 1983.

(339) McCarthy's REMINISCENCES ON THE HISTORY OF TIME SHARING.

There are a number of accounts of this narrative. First, from the first-person source:

McCarthy, John. " Reminiscences on the History of Time Sharing". 1983.

McCarthy, John. " Memorandum to P.M. Morse proposing Time Sharing". Quoted in A Century of Electrical Engineering and Computer Science at MIT. K. Wildes, MIT Press 1985, p243.

Another source is Judy O'Neill's doctoral dissertation (the University of Minnesota, 1992) on the birth of timesharing. Other sources concerning the initial invention of timesharing include the various CBI Oral History Interviews, notably those of McCarthy, Jack Dennis and Fernando Corbato. Transforming Computer Technology studies this topic exhaustively, as does Levy 1984 and Tracy Kidder's The Soul of a New Machine 1981.

Atsushi Akera's Calculating a Natural World (2008) contains essays on the Cold War and education, as well as specifically on "Discipline and service: research in computer time sharing at MIT and the University of Michigan".

(340) McCarthy 1959, 1983, Levy, Hackers.

(341) Levy, Hackers, 67-68.

(342) Campbell-Kelly and Aspray 1996, 209.

(343) Levy 1984, p25.

(344) Levy 1984, p27. However, as early as 1956, freshmen at Carnegie Tech could take programming courses. While Levy's orientation is toward hacking as the genesis of the computer revolution, sometimes revolutionary work is done by people taking regular coursework.

(345) McCarthy, CBI OH interview 156.

(346) Patrick Winston CBI OB Interview 196, 1990; also Ceruzzi 2003.

(347) Lawrence G. Roberts, CBI OH Interview 159, 1989.

(348) Simon also worked with economists, but that cohort did not take an interest in AI until much later.

(349) Bellman was a distinguished mathematician who wrote fundamental texts in dynamic programming and matrix analysis.

(350) The RAND Corporation: the First Fifteen Years, 1963.

(351) Marvin L. Minsky, Charles Babbage Institute OH Interview #179, 1989.

(352) Taube 1961, p9. Taube's book, inspired by a [Science](#) article by Norbert Wiener, seems to have been rather widely read, but is a poorly written compendium of most of the anti-computing arguments taken together.

(353) Turing 1950, Turing had pointed out that this is, practically speaking, not true for a rather mundane reason: "Machines take me by surprise with great frequency. This is largely because I do

not do sufficient calculation to decide what to expect them to do...”

(354) Turing 1950, 27.

(355) Wiener 1960, p1357. Norbert Wiener; Science May 6, 1960: “ Some moral and technical consequences of automation”; 1355-1359.

Wiener had been warning of the possibility of computers and robotics replacing humans and thus causing a depression which “will ruin many industries- possibly even the industries which have taken advantage of the new potentialities.” The Human Use of Human Beings: Cybernetics and Society, 1950, 189.

(356) HAS indicated that this is what Edward Feigenbaum remembers; MOML p206.

(357) Simon, H.A. and A. Newell. “ Heuristic Problem Solving: The Next Advance in Operations Research.” Operations Research 6 (January-February 1958): 1-10.

(358) Bellman responded angrily in another issue of the journal Operations Research, May-June 1958, 6, no 3.

(359) McCorduck 1979, 186.

(360) Spring 1966: 258-264. Simon, “A Computer for Everyman” in The American Scholar, Spring 1966: 258-264.

(361) “Machines with Experience”, Time Magazine, volume 72, n23 (December 8, 1958): p73.

(362) Taube 1961, 26.

(363) Pfeiffer, The thinking machine: everyman’s introduction to the world of electronic devices, 143.

(364) Booth 1965, 123.

(365) Quoted in Waldrop 1987, 63.

(366) Taube 1961, 26.

(367) Weaver's original article 'translation' in Locke, William Nash and A. Donald Booth, eds. Machine translation of languages; fourteen essays. [Cambridge] Published jointly by Technology Press of the Massachusetts Institute of Technology and Wiley, New York. 1955. See also Yehoshua Bar-Hillel, "The present status of automatic translation of languages". Advances in computers. New York, Academic Press, 1960, and Automatic Machine Translation by Oettinger, 1960.

(368) Booth 1965, 121.

(369) Oettinger, Anthony G. Automatic Language Translation: Lexical and technical aspects with particular reference to Russian. Cambridge, MA: Harvard University Press. 1960.

(370) Taube 1961, 27.

(371) Hutchins 1995.

(372) Pfeiffer, The thinking machine: everyman's introduction to the world of electronic devices, 148.

(373) Slocum 1984, 3.

(374) Slocum 1987 article.

(375) McCorduck 1979, 176.

(376) Pfeiffer, The thinking machine: everyman's introduction to the world of electronic devices, 130.

(377) Simon 1991, 272.

(378) August 1955: TR 55 transistor radio; The first popular portable radio. Newsweek 10 25 2004, p88.

(379) Hafner and Markoff, 38. Shurkin p304-306.

(380) The source for most of the materials above regarding NASA, NACA, Sputnik and their brethren is the Aeronautics and Astronautics Chronology, compiled by NASA Historian Eugene M. Emme, and found on the NASA website. McDougall's The Heavens and the Earth also provides a detailed narrative.

(381) February 7 58: The Advanced Research Projects Agency (ARPA) was established by the DOD, and Roy W. Johnson, a vice president of General Electric Co., was appointed by Secretary of Defense McElroy as its Director. ARPA was placed in charge of the Nation's outer space program. Killian, Sputnik, Scientists, and Eisenhower: A Memoir of the First Special Assistant to the President for Science and Technology.

(382) Ibid.

(383) Dickson, Think Tanks, p11.

(384) Licklider 1988, p219; Licklider, CBI OH Interview 150, 1988.

(385) Jack Ruina, CBI OH 163, 1989.